



INFERENCE BELIEF AND INTERPRETATION IN SCIENCE

AVIJIT LAHIRI



INFERENCE BELIEF
AND
INTERPRETATION IN SCIENCE

Avijit Lahiri

Inference Belief and Interpretation in Science

(an e-book)

All rights reserved.

by **Avijit Lahiri**, author and publisher.

252 Laketown, Block A, Kolkata 700089, India.

e-mail: avijit.lahiri.al@gmail.com

website: physicsandmore.net // physics-plus-more.blogspot.com

[May, 2019; January 2020]

June 2020

Contents

1	Introduction	1
	Opening words	1
	Inference and reasoning: the intrinsic and the extrinsic	3
	Perception and cognition	7
	Concept: the basic entity of thought	9
	The substratum of cognition	10
	Logical leaps in inference	11
	What science aims at	11
	The logical and the extra-logical: inferences and beliefs	13
	Belief-emotion-feeling: the unacknowledged trinity	14
	The received view — and beyond	15
2	Objectivity	18
	The irreducible gap	20
	The observable and the unobservable	23
	Observation involves interpretation	24
	Observations are partial: multiple layers of description	27
	The existence of objects, and their relatedness to one another	30
	Theories as condensed descriptions of inter-relations	31
	All observations and theories are contextual	32
	Theoretical concepts: greater and greater depths in a layered description	35
	The big question: how do theories correspond to reality?	37
	Summary: the issues of relevance	40

CONTENTS

3	The emergence of theories: how are theories constructed and accepted?	48
	The essential tension: the objective and the subjective	49
	Science on a pedestal	50
	Logic and reality	51
	Michael Polanyi: roots of personal knowledge	53
	Thomas Kuhn: the paradigm shift	60
	On the irrationality of the substratum	62
4	Inductive inference	64
	The problem view and the process view	66
	Inductive inference: contextual, ambiguous, and non-deductive	72
	Induction: definition, features, and taxonomy	74
	Can induction constitute a logic of confirmation?	81
	Induction: questions and issues	84
	Unconscious cognition	85
	Induction: the role of heuristics	88
	Induction: beliefs emotions and affects	91
	Inductive inference: summary	97
5	The cognitive unconscious	101
	The unconscious as the arena of complex cognitive processes	103
	The unconscious versus the conscious	104
	Unconscious cognition: the role of emotions	105
	The unconscious detection of similarity	105
	The 'rationality' of the unconscious	106
	Unconscious cognition: how 'hot' is it?	109
6	The process of inference: beliefs and emotions	110
	Beliefs and belief systems	110
	Belief, knowledge, and inference	112
	Beliefs and emotions	116
	Heuristics in the inferential process: the dual role of beliefs	121
	Deterministic but unpredictable	126

CONTENTS

Heuristics as inter-subjective and person-specific rules	127
7 Reasoning and rationality: the intrinsic and the extrinsic	130
The rationality issue: a brief overview	131
Rationality: misplaced notions	137
The dual-process theory	141
Complexity	147
The unpredictable	147
The essential role of the context	149
The rules of inference-making	151
Guessing at the inner mechanisms of a complex system	154
8 Abduction: theory in emergence	156
The generation of hypotheses: an adventure into the mystery world of psychology	158
Abduction and novelty	159
Abduction: the naturalist point of view	163
Abduction: the crucial link in science	164
The process of abduction: exploration of a conceptual space	164
Amplifying mechanisms, instabilities, and parallel explorations	166
Abduction: the setting up of new correlations	170
The reasoning individual: an atom in a cosmos, and a product of evolution . . .	171
The parable of the two children	172
Analogy: the great organizing principle	172
Further principles: simplicity and elegance	173
Abduction: the inspiration of science	174
9 Suming up: science as an interpretation of the world	175
Scientific theories: the basic issues of relevance	176
Realism or anti-realism?	178
Theories: from the immediate to the remote	181
The two facets of truth	182
Hypothesis and theory: the individual and the scientific community	189
The aim and method of science: the question of values	192

CONTENTS

Science as a telescope: the ‘objective’ and the ‘eyepiece’	194
Approximations in science: the metaphor of convergent and divergent series . .	196
Asymptotic series and singular approximations	196
Singular reduction from one theory to another	199
The crossing of borders: scientific revolutions	201
Normal science: conceptual restructuring at all scales	202
The complex relation between theories: emergent phenomena	203
Contextuality in science: summing up	206
Theory choice: the problem of ‘underdetermination’	207
Scientific progress: socially determined or socially conditioned?	211
Evolutionary psychology and cultural inheritance	214
The rationality of science: values in troubled times	216
10 Summary and concluding words	221

Chapter 1

Introduction

Should you be interested in this book?

Opening words

Science is commonly perceived to be the last word in logic and objectivity, where the latter has two aspects to it — one of being person-independent and the other of being a true description of the workings of Nature. In this little book of mine I will tell you, as best as I can, of the possible presence of extra-logical components in the reasoning process of individuals and groups of individuals, making special reference to the context of scientific inquiry. And I'll also indicate the possibility of a skewed fit between what science tells us about nature, and Nature itself. In the process, we will get to understand how all this can constitute a strength rather than a deficiency in the way science inquires into nature.

To what extent do people conform to standards of logic and rationality in everyday reasoning and in scientific activity? How do the extra-logical components woven into the reasoning process of individuals lend color to the way they solve problems relating to workings of nature? To what extent does inter-personal and group interactions succeed in effacing the subjective component involved in individual and collective reasoning and yield objective knowledge, transformed from the belief-laden knowledge of individuals

and groups to knowledge gained by humankind?

These and related questions have repeatedly been addressed by authors for a long while now. I will collect here a few of the basic ideas that have come up in the course of studies into these important and interesting issues. In the process, I hope to tell you, in outline, how and to what extent scientific inquiry can possibly be conditioned by the mindset of people.

In recent decades, philosophy has taken a naturalistic turn. Briefly stated, naturalism looks at things and processes as they actually are, without looking at these from too abstract and analytical a perspective — a tendency so common in philosophy. In this, the naturalist point of view follows the point of view of science itself (see, for background, [48], [43], chapter 1). The issues relating to mechanisms underlying the inferential processes in men, and to human rationality at large, are now being addressed from a naturalist point of view. The branch of inquiry dealing with these and many other related questions on mechanisms of human thinking that has developed over the last fifty years or so is known as cognitive science. In this book, I will have occasion to refer to the naturalist point of view and to ideas in cognitive science, as someone who has been keenly interested in these in the context of philosophy of science, someone with a background in physics.

My job in this book will be to share with my readers a point of view that has a good degree of contact with current literature on a broad area in cognitive science and philosophy of science. I will explain a number of basic concepts to set the tone of this book and then go on to propose a framework that raises deep questions on the received view of science, based on cognitive aspects of our reasoning process. In the process, I will venture to place before you my own interpretation and speculation of how things work in individual inferential processes that find expression in scientific exploration by individuals and by scientific communities at large, again taking care that these do not conflict with the body of opinions and beliefs shared by experts in the relevant areas of discourse.

I will do my best to build up a picture that hangs together, is not inconsistent with what experts have found and written on the issues involved, and is, one hopes, relevant and

interesting.

Here are a few words on how the text is displayed in this book.

As you can see, this book runs in two courses — one written in larger font that constitutes the main text, and the other in smaller font that makes up a sub-text. The latter is for the purpose of additional explanation and clarification, along with references to literature. There will not be too many of the references, and most of those will be to books and monographs. You will find only few references to journal articles, since I do not want to burden the presentation with technicalities. You can read the text and the sub-text *any way you like*.

This said, let us move ahead. I will, in the remainder of this opening chapter, take you to a brief tour of some of the key ideas that will define the content of this book.

Inference and reasoning: the intrinsic and the extrinsic

The first basic idea I wish to highlight is that of the role of inductive inference in reasoning in general, and in scientific inquiry in particular. Induction and deduction are commonly taken to be two distinct modes of reasoning where, in either of the two, one starts from one or more premises and then draws a conclusion. This act of starting from a set of premises and reaching at some conclusion I will call inference. In this act of inference, a person may, depending on circumstances, make some use of rules that may have something to do with logic. Now, this may seem to be a pretty roundabout and tentative way of putting things but, in this book, you will soon get used to such roundabout statements, and qualifications (with or without the use of parentheses), because most of the time we will be treading on murky ground.

Going back to inference and the use of rules, we will say that the act of inference conforms to the description of *reasoning* whenever some set of rules of general validity are involved, applied sequentially. Now, the concept of rules is a tricky one since there are two distinct contexts in which one can talk of rules — a distinction that, at times, goes unnoticed. As a person gets into an act of inference, she may be following, in her own

way, a set of rules of her own where, to make matters complicated (and interesting too!), many of the rules may be unknown even to herself because they operate unconsciously. These we will call intrinsic rules for the sake of easy reference. On the other hand, somebody, say, a cognitive psychologist, looking at her inferential act may try to fathom out what she (the subject) is doing, by trying to relate the relevant premises and the conclusion by means of a set of extrinsic rules, rules called in use by her, the psychologist. How are these two sets of rules related?

In a sense, this is a fundamental question not only in the area of inquiry we are looking at, but in all areas of human inquiry, being central to the concept of science itself. I, like many others, have always found it fascinating. This is a theme we will time and again come back to. As I see it, it is a basic question raised by naturalism. Think of the following scenario: a sequence of numbers is being taken out from a machine — never mind how or by whom — and the numbers are being displayed to you. You are trying to fathom out the rule underlying the sequence in which the numbers appear. You make a guess and find that the guess is working — only up to a point. As more numbers come out, you make a fresh guess and are again rewarded, but once again you find that the sequence is more inscrutable than what your guess tells you. The rules you are guessing at and comparing with the sequence at any point of time are the extrinsic ones, which don't quite match with the 'rules', if any, by which the sequence is made up — the intrinsic ones. For all we know, the sequence may be of a random nature, with a number of regularities built into it that make you arrive at guesses and let you be hopeful of eventually arriving at the 'correct' intrinsic rules. Nevertheless, the intrinsic and the extrinsic remain inexorably distinct. The business of science is no different. Here, it is Nature that causes a multiple sequence of events for mankind to guess at the rules underlying the occurrence of those. Naturalism firmly distinguishes between the guesses and the actual workings of nature, and does not entertain abstract notions of whether and how the former may approach the latter.

Cognitive science tries to guess at the intrinsic rules, if any, involved in the inferential reasoning process of individuals by comparing it with a set of extrinsic rules, and the

starting point in this endeavor is provided by what may be called the *rules of logic*. Here the term ‘logic’ really refers to deductive logic, because people also make frequent use of the term inductive logic, thereby referring to a distinct type of inferential activity. It makes more sense, however, to speak of inductive inference rather than of inductive logic, and to use the terms ‘deductive logic’ and ‘logic’ interchangeably.

At a deeper level, the concept of ‘intrinsic rules’ of nature (or, in a different context, of intrinsic rules operating in inferential processes), taken literally, is somewhat a misplaced one. What we think of as a set of intrinsic rules may be just a set of interrelations, or correlations, among things in nature (or among cognitive factors in the process of making of inferences). Science aims at understanding or ‘reproducing’ (in its models of the world or of the human mind) these correlations, but in its own interpretational terms; what is ‘intrinsic’ to the world or to a cognitive mind, is sought to be replicated ‘externally’, by means of ‘laws’ and ‘rules’. Inherent to the idea of laws and rules is the supposition of some design (not necessarily, though, by some superhuman ‘mind’). The correlations in nature, or among cognitive factors in the human mind, need not be based on any ultimate design that is there for us to decipher.

In more precise terms, one has to distinguish — as far as extrinsic rules are concerned — between various classes of logic like, for instance, propositional logic, predicate logic, doxastic logic, and deontic logic or, more generally, instances of *modal* logic, some of which may partially overlap with one another in connotation. However, all of these are of the nature of deductive logic since they relate to ways of inference involving the operation of unequivocally defined inferential rules on premises whose meanings in some bigger context are not of direct relevance. As for now, all we need to know is that deduction and induction are two types of inferential activity that can, at least provisionally, be distinguished from each other. In this scheme of things, reasoning can be described as an inferential activity that has some correlation with rule-based inference, where the rules, generally speaking, are not specific to individuals, and have some degree of universality such as the rules of deductive logic. As you see, there is no sharp distinction between reasoning and inference, which is only to be expected of the issues in human

cognition we will be discussing in this book.

The various instances of modal logic are designed to capture different aspects of thought and reasoning, as they operate in real life, but from the logical and computational point of view. These find applications in a wide variety of disciplines and are, in particular, of vital relevance in artificial intelligence. An overview of modal logic can be found in [7].

Indeed, as we will see, the distinction between deductive and inductive inferences is also not so sharp. Once again, if this sounds confusing, you better get used to it because cognitive psychology and some other parts of cognitive science are way apart from physics or chemistry, and require a different mindset. Be warned that our job here is, primarily, to get exposed to a number of ideas rather than to try to define those too precisely and too prematurely. In this, we will be following the spirit of a critique to the conventional image of science itself, an image that is increasingly being brought into question in recent decades.

Aspects of the conventional image of science, one that has lingered in philosophical and logical accounts of science, as also in popular perception, have been collectively termed the 'legend' by Philip Kitcher ([71], chapter 1). Kitcher is a notable name in contemporary philosophy of science in his broad, perceptive, and nuanced approach to issues in science, including those relating to *social* aspects of science in today's world.

In the remainder of this chapter I will introduce a few other terms that will come up repeatedly in this book (as these do in our everyday discourse as well) in connection with human mental activity, namely, perception, thinking, and cognition, of which, the last term is a bit more specialized compared to the other two. And again, I will be no more specific than telling you that I will use these terms in the sense of *common usage*.

Perception and cognition

We perceive something when a set of stimuli act to influence our senses (typically, the sense organs) so as to create an impression of the source(s) where the stimuli originate from. Frequently, perception involves an internal processing of information whereby our mind comes to form an idea of that source of information (such as, the appearance of an external object in the case of visual perception) conforming, in some sense, to the nature of the source itself (thus, we perceive a table as a table, though there exists serious debate, not all vacuous, as to how much our perceived table is the same thing as the table itself).

There has been protracted debate as to whether our senses give us an unblemished picture of the world. Each and every act of perception is, in a very basic sense, an act of interpretation in terms of past experience. “No one now seriously believes that the mind is a clean slate upon which the senses inscribe their record of the world around us: that we take delivery of the evidence of the senses as we take delivery of the post”, writes Peter Medawar, a pioneer in the field of immunology and in the science of tissue transplantation, and philosopher of science of great clarity of view [88]. Medawar goes on to quote Nietzsche: “Everything that reaches consciousness is utterly and completely adjusted, simplified, schematized, interpreted,....”, and Kant as well: “experience is itself a species of knowledge which involves understanding....”.

Here is one more instance of how perception involves complex psychological processing of information received from the world:

“Perceptual processes were at one time believed to be lower-level functioning, both because they are accomplished without our conscious control and because even animals can do such things as recognize patterns and learn spatial layouts.... . Perceptual processes were contrasted with the higher mental processes, such as problem solving, logical thinking, and decision making, which are accomplished consciously, and are much less obvious in animals. There is, however, a basic problem with labeling perceptual processes as lower-level: The ability to recognize patterns and direct our attention involves very elaborate cognitive computations, which are heav-

ily influenced by top-down processes that depend on ones knowledge and interpretation of situations. Thus, it is very difficult, if not impossible, to separate lower and higher forms of cognition.” ([125], p181).

The idea that all our perception of nature is, fundamentally, an interpretation, will constitute a recurring theme in this book.

Incidentally, we often use the term ‘perception’ in a deeper and broader sense (*my perception of the current political situation differs from yours*). This variation in senses attached to the same term is very common when we try to say something of our mental world, since the mind is an infinitely complex and flexible whole, where it is almost impossible (but often necessary!) to meaningfully distinguish between parts of it or between mental processes of various descriptions.

When we say that perception involves some processing of information, we are close to the sense carried by the term ‘cognition’, because cognition is commonly understood to be that mental activity where there takes place some processing of thoughts and ideas, tied to some purpose. What is important and interesting in this context is that some (or most) of the processing may be of an unconscious nature, which we ourselves may not be aware of. Cognition, in other words, refers to a broad class of mental activity that includes inference making and reasoning in which, generally speaking, there is a conscious component. In this, inference making and reasoning are special types of cognitive activity where there is a relatively greater role of a ‘purpose’ or a ‘goal’ guiding the course of that activity though, again, the purpose or goal may be, to a large extent, hidden from our own conscious awareness.

While perception and cognition are broader in scope than inference making and reasoning, ‘thinking’ is, likewise a broad and all-inclusive term with unconscious and conscious processes involved in it. Of, course, like everything else, the term ‘thinking’ is used in senses often widely differing from one another (*‘what were you doing?’*,— *‘oh, just thinking’*; or, *‘I am thinking hard how best to entertain our guests’*); but, well, what can one do?

Concept: the basic entity of thought

One other term will be especially relevant for our purpose in this book: *concept*. What is a concept? As I see it, a concept is some kind of organized thought about some object or set of objects (or about some other concept(s)) that tells us something about the object and that is made use of in any context involving that object. For instance, my concept of a table is involved when I ask somebody to move our dining table to some other place in the room (the concept tells me that the table is something that can be moved, by the application of some particular kind of effort, without getting altered in any appreciable way) or when I am engaged in a philosophical debate, claiming that the table is not entirely what it seems to be (the concept tells me that there are aspects to the table not perceived by me). The term concept holds a huge significance since it is by means of concepts that we understand and organize the world around us and, what is more, act in (and upon) it. Concept, in brief, is the thought entity by means of which we make sense of our world and of our life. A concept commonly involves a central entity (the table) associated with numerous other entities stored in the mind, all these woven into a single whole. Concepts, in turn, are associated with one another at various levels so as to make up more complex structures like ideas, beliefs, and items of knowledge. I emphasize the term ‘association’ here: association between thought elements in the formation of concepts, and association between concepts in the formation of more complex forms of thought.

In summary, I have introduced to you the terms ‘reasoning’ and ‘inference’, ‘deductive’ and ‘inductive’ types of inference, and ‘logic’. I have also mentioned terms that will hover in the background throughout the discourse presented in this book — ‘thinking’ (or ‘thought’), ‘cognition’, and ‘perception’. And I have told you what the term ‘concept’ will be meant to stand for. This will constitute the starting point to the next phase of our journey where we will dwell upon these ideas at greater length, and upon many other related concepts and ideas.

The substratum of cognition

Let me go back to the point where I digressed from, namely, inductive inference. While the basic idea of inductive inference will be one constituent of the central theme that will define this book, the other constituent ideas of dominating relevance will relate to the *substratum* of mental activities that determine and guide the process of inductive inference — activities that belong to a substratum in the sense that these have traditionally been ignored in the understanding and analysis of human cognition, and also in the sense that it involves a multitude of mental processes that the individual (or group of individuals) engaged in making the inference is often not aware of. I refer to the substratum of the *cognitive unconscious*.

This idea of a substratum is of relevance in every sphere and at every level of human inquiry. It is relevant, in the sense indicated above, in the context of the way we perceive the process of cognition. More generally, it is relevant in the way science inquires about nature. Any such inquiry is, by the very nature of things, limited to some specific domain or other, and there is a substratum that is either ignored as being not of relevance or whose existence is not known or suspected. But, unknown to us, the substratum holds the key to many things. For instance, in the domain of electricity and magnetism, the substratum, at one stage of discourse (before the introduction of unifying ideas by Faraday and Maxwell), was constituted by the wave nature of electrical and magnetic disturbances, which held the key to a vast range of phenomena involving the joint variation of electrical and magnetic field strengths. At a different level, within the context of the classical electromagnetic theory, the substratum relates to the quantum nature of the electrical and magnetic fields. This is another theme that will be repeatedly encountered in this book: every inquiry is limited to some specific context or other that may or may not be explicitly acknowledged, underneath which lurks a substratum — one that is either not known to us or about which we are not aware. The entire perspective of the inquiry changes as the substratum is recognized and taken into account, whereby a new context is set.

Logical leaps in inference

Continuing with the question of inductive inference, I will, in this book, devote quite a few words to induction and deduction, broadly distinguishing between the two (I will outline to you the sense in which these two are distinct and also the one in which they overlap), emphasizing time and again that inductive inference is a process involving *logical leaps*, i.e., gaps in the sequential application of rules of a more or less general validity ones of well-defined nature. This is an all-important and pervasive feature of human thinking that begins at the level of the most trivial mental activity of the child and of the grown-up (and, of course, the trivial is only seemingly so) and ending at the level of the most intricate scientific reasoning of the individual and the scientific community. The logical gaps entail, inevitably, the necessary role of *choice* between alternatives, where a choice is a selection that is not dictated by rules independent of the specific context in which it (the choice) is made. This will again require that we address the question as to what the choice really and precisely is and how it is actually exercised, and will again make us confront the question of conscious and unconscious factors in determining the course of human inferential activity. Incidentally, whenever we exercise a choice, we actually make a *decision*. The exercise of choice and the making of decisions, these are ubiquitous at all levels of human activity though we are not always aware of these. The surgeon makes a choice and saves the patient's life; or, the general exercises a choice and makes a strategic retreat it is only such momentous events that engage our attention. But, unknown to us, choices and decisions continue to be made incessantly, and this is what makes our very existence possible.

What science aims at

Science is done by the individual scientist on the one hand, and by the scientific community on the other. I will not try to define who the individual scientist is, or what differentiates her from her fellow men. And the scientific community will remain similarly undefined, like so many other things in this book, because explicit definition does not always add to understanding.

CHAPTER 1. INTRODUCTION

In quite a considerable number of issues addressed in this book, I will rely on what I think is the *common* ground between me and you, and will not try too hard to make things 'evident' — all I want to do is to share, and not to impart. Sharing begins on common ground and ends up expanding it.

The scientific thinking of the individual is generated and processed in the labyrinthine innards of the human mind and, as such, is conditioned by and stamped with individual idiosyncrasies. At the same time, the initiation of the thought process and the final product of it have to relate, at least in some sense, to what the *scientific community* thinks and does. It is the community that ultimately decides whether and when the intellectual product offered by the individual scientist gets to be integrated into mankind's storehouse of knowledge. Here we will have to confront the question as to how far the socialization of individual thought leads to an 'objective' view of nature, as it is commonly supposed to.

And this will open up another vital issue: what is science for? Is science a means to solve, and to keep on solving, problems faced by men and groups of men as they move along in the business of their life, a means to continually ensure and improve upon their survival and existence? In short, is science a means and a strategy to solve problems faced by men and women in this world? Problems relating to disease, hunger, and the innumerable other aspects of living? Or, is science a strategy to probe into the workings of nature, into the hidden secrets and mysteries that make nature what it is? How much of a practical necessity is science and how much of an aspiration to Truth? Can it be that the two are so intimately related that it is futile to seek an answer to this question either way? Or, can it be that by recognizing the distinctive natures of these two aspects to science we can achieve clarity in addressing and answering some other questions of concern to us? I don't really know, but we will see

Whatever be the purpose that the individual or the scientific community may want fulfilled by doing and practising science, the effort to do science involves immensely complex motivations, aspirations, tensions, and intellectual resources. Some of these are within the individual, and some within the community, but the two merge unto

each other in unfathomable ways. In this book, my attention will be focused once onto the individual and once onto the community, but always with the implied admission that the individual is a microcosm of the community and the community carries all the contrariness of the individual. Science begins in the mind of the individual and ends up in the knowledge, belief, and practice of the community. And again, it begins in the concerns of the community, and ends up blossoming in the mind of the individual.

By the term ‘community’ I will primarily mean the scientific community within which the individual scientist works; at times it will implicitly stand for smaller groups of workers in the immediate periphery of the working scientist with whom she collaborates and shares her thoughts; and at others it will mean the social environment at large that nourishes and nurtures her and, at the same time, makes her a part of a complex process.

The logical and the extra-logical: inferences and beliefs

And now I feel I have set the tone of this book: Within the broader question of what we want science to do for us — how much we do science with a view to understanding the mysteries of nature and how much to acting upon nature to achieve certain ends — this book will take up the question of how science is done by the individual situated within a broader context — the context provided by the scientific community on the one hand, and by the broader social and cultural environment on the other. In the pursuit of both these questions, we will examine the role that logic plays in the scientific activity of the individual and of the scientific community, and in shaping our concept of nature, and we will look at the nature and extent of extra-logical aspects in the inferential processes involved in most of our everyday activity and in science. These extra-logical aspects relate to the logical leaps and the choices that are forced upon us in facing and confronting the world around us, in confronting Nature if you wish, where both conscious and unconscious factors contribute to the process of taking a leap and making a choice.

In this book, I will adopt the view that doing science is continuous with our everyday, mundane inferential activity, while being an exceptionally special and focused instance of the latter.

In this, the two recurrent themes that will be thrust upon your attention will be: *inductive inference*, and *belief*. Belief is commonly, and quite justifiably, held to be the harbinger of superstition, pseudo-science, and bigotry. At the same time, belief is also commonly held to be the antithesis of logical inquiry. But here *I will try to present a different perception of the role of belief*.

Belief-emotion-feeling: the unacknowledged trinity

Belief is that vast marshy ground that supports both superstition and bigotry on the one hand, and the most astounding inferential feat on the other. And it is exquisitely difficult to disentangle these two aspects of belief from each other. Part of the difficulty lies in the fact that there has not been much relevant work specifically on belief in cognitive psychology, while not an inconsiderable part also relates to the nature of belief itself. Belief is not knowledge, but it provides that substratum on which knowledge rests and unto which knowledge merges continuously, quite seamlessly, and with little differentiating shade. In inductive inference, belief provides the springboard for logical leaps, the compass needle that, rightly or wrongly, lets us select, choose, and decide among alternatives for which pure logic supplies no clue. And in this, *belief is aided and abetted by emotion and feeling*. Belief, emotion, and feeling, these make up the triumvirate, operating mostly in the unconscious world of men, that I believe to have a great but subtle role in guiding and shaping their scientific quest, a quest commonly perceived to be the pinnacle of conscious, logical inquiry on the part of humankind. In this book I will tell you if there is ground for this *belief* of mine.

The received view — and beyond

Indeed, this book will, in its own way, pose a critique to the commonly held image or, in other words, the ‘received view’, of science. As I have stated above, recent decades have witnessed a questioning, from various different quarters, of the traditional image of science and, as a matter of course, of science itself — of what science has been doing to our world. I will draw from some of these other critiques. I will, in particular, look into the question of how objective a view of nature science provides us with and examine how fragile that view is. This, along with a number of other similarly important issues, will constitute an implied re-examination of the commonly held perception of science, if not of science itself.

The origin of the ‘received view’ is diffuse, since this view resulted from diverse accounts of science and the scientific method given by scientists and philosophers, and other men of eminence, mostly belonging to the western intellectual tradition. Roger Bacon and Francis Bacon are considered as initiators in the building up of the perception of what should count as the method of science. The British empiricists contributed greatly to the further consolidation of the view, to which a logical foundation was added by the logical positivists and logical empiricists of the last century. Many of these trends within the broad umbrella of the received view were, of course, remarkably acute and prolific, and were sharp enough to lay the foundation of a critical examination of the received view itself. Karl Popper, while an outspoken critic of logical positivism, was among the last great architects of the received view, and attempted the formulation of a sharp criterion for the demarcation between what is scientific and what is not. The *demarcation problem* subsequently lost its initial promise of providing a precise definition of what science is supposed to be. Ironically, the Marxist trend in social movements contributed to the perception of science as a disembodied intellectual process, realizing the immutable laws of dialectics. It added to the received view the tenet of social determination (but not social construction) of the course of science that eventually turned out to be a simplistic and suffocating point of view while, at the same time, retaining the potential to pose a substantial critique to science in the present day world.

The counter-current to the received view began with a naturalistic turn to the philosophy of science whose origins were, however, yet more diffuse. Results of the naturalist turn in the philosophy of science were augmented by trends in cognitive psychology that emerged as the behaviorist approach in psychology lost momentum. Naturalism, moreover, is itself not a sharply defined point of view, and various naturalist trends are often devoid of a common and strong family resemblance. Michael Polanyi was among the early critics of the logical positivist and logical empiricist view of science as the repository of impersonal, logical, and objective knowledge, progressing cumulatively to a ‘true’ picture of nature and its inner mechanisms. N.R. Hanson, the ‘Flying Professor’, raised doubts against the idea of an ‘objective’ confirmation of scientific theories by observed data, pointing out that observation itself was theory-laden and, drawing upon the pragmatist tradition in the philosophy of science, contributed in no small measure to the inauguration of the cognitive-naturalist era. And then came the final onslaught by ideas unleashed by Thomas Kuhn [76], aided quite considerably by a criticism of the received view of science by Paul Feyerabend [38].

Philip Kitcher [71] has drawn a distinction between ‘legend bashers’ and ‘science bashers’. Among the former are those critics of the received view of science who aim at transcending its limiting horizon and arriving at some position beyond that horizon. Kitcher is himself a critic who, nevertheless, aspires to identify a ‘legacy’ to the legend that has been built around the received view.

Science bashers, on the other hand, will not be referred to in this book beyond making the remark that science bashing is the necessary obverse of science fetishism which completely ignores deep questions relating to the idea, seemingly ingrained in science, of taking control over Nature. Science, indeed, is in trouble in today's world since it has become, to all intents and purpose, synonymous with control and power.

These few remarks are, of course, too sweeping to be taken seriously. I include these as being indicative of the spirit in which this book would like to view the commonly held perception of science. Precisely because of the sweeping and personal nature of the appraisal presented in these remarks, I do not include a great many references to substantiate what I say. However, I consider [71] and [72] as general references that you may find useful, and illuminating too.

CHAPTER 1. INTRODUCTION

The themes raised in this chapter will run through the course of this book.

I will move on from here.

Chapter 2

Objectivity

How truthful is science?

Or

Does science REALLY describe the workings of Nature?

This chapter will be concerned, principally, with *scientific realism*.

I will examine, to the best of my understanding, how objective our conception of Nature, gained through Science, is. Now, the term 'objective' may have different connotations in different contexts. One possible connotation is: person-independence. We would prefer the findings of science to be independent of the idiosyncrasies and the mindsets of individuals, as also of those, if any, of specific communities and cultures. In short, we would prefer science to be a product of the *whole of mankind*. At this point, I will not question as to how far this ideal aspiration is achievable, either in practice or in principle. I will provisionally assume that, somehow, science appears as a product of mankind without any birthmark resulting from its origins in individuals and particular scientific communities. In that case, the second connotation of the word 'objective' will demand attention and examination: how accurately does Science tell us what Nature

actually is, and what its inner workings are?

The question of impersonality of science will be addressed in subsequent chapters (mostly in chapters 3, 5, and 6), where we will get to face the contrariness inherent in the process of science: science purports to be knowledge possessed by mankind, but it develops in the minds of individuals — minds deeply riddled with beliefs of a personal nature alongside of ones less personal and more objective.

This, in brief, is the question that scientific realism addresses.

There exists quite a vast literature on scientific realism and critical appraisals of its tenets from various different points of view. This is a subject of hard core philosophy which it is way beyond me even to think of outlining or summarizing, let alone attempting an in-depth review on. Scientific realism as a point of view, or as a philosophical position, has had its beginnings in antiquity, and continues till date to spawn heavily polemical literature of various shades at a rate that shows no signs of abetting. The polemics proliferates into ever-expanding areas, in pace with scientific theories themselves.

Further reflections on realism and anti-realism, two oppositely oriented points of view, are to be found [here](#) in chapter 9.

My aim in this chapter is, as a matter of fact, quite modest, one which neither the realist nor the anti-realist will, perhaps, consider worthy of objection in any strong measure. In the end, however, I will likely be judged closer to the realist position, and the anti-realist may not be happy. The realist also may not feel comfortable in my company, but I will not worry too much on questions of who thinks what, and will now get down to what I want to say to you on how and to what extent science describes nature and explains its workings.

Recall that I want this book to make you aware, if you are already not so, that the conception of nature that science builds up, is rather riddled with logical gaps or, in

other words, that science is not as logical a business as it is commonly made out to be, being in the nature of an interpretation of the world. It is the existence of the irreducible gap between the interpretation and that which is interpreted, that, on the face of it, speaks against realism. At the same time, this book does not want to give you the impression that science is not much more than a convenient and clever exercise, with little engagement with reality, where it is incapable of making authentic statements about nature, which is where it will differ from the anti-realist point of view as well: if science constitutes an interpretation of the world, then the world has necessarily to be supposed to exist independently of the interpretation.

The irreducible gap

The very first thing that, I think, is needed, is to stick to the basic fact that our conception of nature, as built up in scientific theories, is distinct from nature itself. I will not adopt the position of stating that it is meaningless to talk of 'nature itself', which is the stance sometimes adopted in critiques of realism where one underlines the fact that the most that we can say in the matter is to the effect that we have innumerable 'sense data' — various sensations on the basis of which we build up all our concepts and theories.

When we press our palm against a table-top, we feel some specific sensation which, when we come to think of it, is all that we have in this particular instance of our conception of a table. To be sure, we can also see the table standing in front of us, but there again, all we can be 'certain' of is our sensation of vision. Even summing up all our sensations about the table, we can never be sure of the existence of the table itself.

How, for instance, can we rule out the possibility of an omnipotent malicious demon creating in us the illusion of all these sensations of a table, which is how Descartes argued his case for radical skepticism? (Likewise, Descartes argued, you cannot rule out being in a state of perpetual dream). This is the position of idealism that Bishop Berkeley espoused. It may seem to be a strange position to adopt for a mind as remarkably penetrating as Berkeleys. Berkeley's arguments were indeed of great cogency, so much so that if one cares to follow those arguments closely, one will be left with little option

but to reach the position to which he leads ([78], chapter 5).

Rene Descartes (1596-1650) is acknowledged as one of the founders of modern western philosophy. His approach was a foundationalist one in that he looked for irreducible foundations of knowledge. He adopted an attitude of fundamental skepticism towards things empirical and accepted the workings of the mind (that part of it of which we are aware) as the one facet of reality we can be certain of, being the fountainhead of rationality. The other facet of reality the mind works on is a God-given one, distinct from the reality our senses tend to lead us to. In this, Descartes' position was antithetical to the one later adopted by the empiricists who accepted the empirical as the ultimate authentic source of knowledge, and were skeptical towards the authenticity of what the mind infers.

George Berkeley (1685-1753) accepted the reality of the sense impressions but did not accept these as being caused indubitably by an independently existing world. Reality, in other words, consisted only of sense impressions or ideas, and the mind that perceives those ideas. Our conception of an independently existing real world is possibly some kind of a trick that the ideas play on us.

This is one instance of a general problem that we will encounter again and again in our discourse. It is indeed not possible to come up with a purely logical argument taking us from our sensations to the world 'out there', since there does exist an irreducible gap between the two. The realist's argument that our sensations must be sensations *of* something and that, that something is simply the world out there, may appear to be quite 'natural', but that still does not make it incontrovertible from the point of view of logic, because one may conjure up alternative, though apparently weird, explanations of our sensations like the one invoking the malicious demon (or a modernized version of it, namely, the so-called 'brain-in-the-vat'; there exist entire semi-religious philosophies saying that the world is nothing but one great illusion). Arguments between realists and people going by these theories of mind-dependent reality — a reality that is a construct of the mind making use of the sense impressions — can be endless as also pointless since the worlds of the two protagonists never meet.

The ‘brain-in-a-vat’ (see, for instance, [99]) is a certain type of thought experiment supporting the position of universal skepticism (akin to Cartesian doubt) towards an empirically sensed reality. Assuming that all our sensations and ideas originate in patterns of neuronal excitation in our brains, it is not necessary to inquire further as to where those excitations come from. A disembodied brain, with its neurons excited in various appropriate patterns can be made to produce exactly the series of sensations that an individual goes through in her life. Once again, there is no sure-shot way of countering such an argument because it is *meant* to be impervious to any such attempt.

The world, indeed, transcends our senses and our concepts (or, to put it differently, our concepts transcend the world), and to reach from the latter to the former is a fundamental and irreducible act of induction that the realist cannot deny, and the more sensibly he comes to terms with this basic fact the wiser he will prove himself to be. It is precisely the gap between nature-in-itself and our conception of nature that implies the possibility of alternative routes in the explanation of our sense data, such as the evil demon, the brain in the vat, or even the world as a big dream or illusion.

The ‘irreducible gap’ between Nature and our conception of it is, however, only a matter of our perception as cognitive agents. More of this later.

It is all very well to accept the realist explanation, of the existence of a real and structured world causing our sense impressions, as ‘natural’, and to dismiss these other explanations as ‘weird’ and contrived, but quite another to use pure logic alone in the act of dismissal because, in the end, all these labels like ‘natural’ or ‘weird’ are extra-logical qualifications that we impose on concepts and theories. Later in this book we will have occasion to look into the matter of the relevance of such qualifications in making inferences in our daily life and in doing science (*the general theory of relativity is a beautiful one; the many-worlds interpretation of quantum theory looks weird to me; he has the knack of coming up with simple and elegant ideas*). All we have to take note of now is that these qualifications do not necessarily make an explanation either compelling or

worthless from the point of view of pure logic. The way we interpret the relation between our sensations and the world around us is a choice that we have to make — one that cannot be done sitting back and looking at formulas in logic. All our life we have to go on applying our mind, look at possible alternatives, and risk making choices — the responsibility for the making of those choices lies with us, and us alone. The choice of accepting the position that all our senses are caused by an independent and structured reality (the senses are also part of that reality) is a huge act of induction that every child goes through without ever being aware of it. And, as she grows up, she renews her commitment to the choice, this time consciously, and continues to get along with it unless, of course, she gets weary of accepting the responsibility of making choices and decisions in a real, uncertain, and troubled world, and swaps position so as to believe now that the world is one big illusion.

The observable and the unobservable

Of course, the acceptance of a real world independent of our minds is not the end but rather just the beginning of scientific realism. While scientific realism has got to say much more than this, it is not, however, one single, neatly defined package. Protagonists differ from one another in what they accept as the defining description of scientific realism. However, pretty much everybody committed to realism has to address the question of a supposed distinction between ‘observable’ and ‘un-observable’ entities making up the world. The table in front of me is one instance of an observable entity while the electron, of whose existence we become aware only indirectly, by certain effects created by it, such as by a voltage pulse in an ionization chamber, is an un-observable entity. There is a philosophical position that tells us that the claims made in scientific theories about the un-observable entities are only convenient means of describing and systematizing various observations, and the entities themselves cannot be assumed to exist in any real sense.

Most realists, however, agree that the electron is as ‘real’ as the table in the dining room, though the means of registering the existence of the two differ from one another.

I will stick by my earlier disclaimer that I will not enter into a detailed analysis of all the various philosophical points of view. In this book, I will state my position on a number of philosophical issues without any appreciable engagement with philosophical literature.

While the statements that I make are not inconsistent with what has been said by specialists in the various issues involved, these will, however, not always be consistent with the position adopted by any particular philosophical camp. In philosophy, issues are discussed and analyzed through debate and discourse, rather than by reference to any independent 'objective' determinant of the validity of this or that position. In looking into what the philosophers say on any particular question or issue, one often finds them to be clustered into groups around leading personalities whose theses serve as nuclei for contending viewpoints. However, I will, at times, cut across the various contending viewpoints and put together the content of this book in what may appear to be a synthetic approach. Broadly speaking, I will be close to the realist point of view, but here again, I will move across various different positions within the realist camp. Indeed, I will be adopting a naturalist orientation while remaining committed to a broadly realist point of view. Rather than burden you with an account of where my philosophical loyalties lie (I do not have strong loyalties to speak of), my aim will be to tell you as clearly as I can, and to the best of my understanding, how and to what extent scientific theories describe nature and where, in this, these bear the stamp of extra-logical inferential leaps.

Observation involves interpretation

To come back to the issue of the table and the electron, I will adopt the position that there is no big difference between the two from the point of view of scientific inquiry. The table top causes our visual sensory organ to be excited in a more or less direct manner while the causal link from the state of the electron to our sensory perception is much more indirect, mediated through a number of intermediate stages. My perception of the table, while apparently a simple and direct act, involves a complex process nevertheless, starting from the excitation of the visual organ, and proceeding through a large

number of transformations and associations in my mind, partly conscious and mostly unconscious. The result of these transformations can, in a sense, be compared with a theory, because a theory is likewise a complex thing even when it appears to be a simple one and because a theory is, like my visual perception, an interpretation of what our sense data, directly or indirectly received, communicate to us. My perception involves, much like a scientific theory, a large number of associations and chains of reasoning, some explicit and some defined only implicitly. An electron causes transformations in the states of systems used for ‘observing’ it, and the actual act of observation is an inference based on some kind of a theory, this time mostly an explicit one, as to how and why such transformations occur. For instance, consider the statement that the voltage pulse in an ionization chamber is caused by ionization of gas molecules by means of electromagnetic interactions. This, of course, is part of a theory, a theory of what electromagnetic interactions are, a theory of how and when such interactions cause the ionization of a gas, and so on. All these remain implied when a scientist says that she has observed an electron.

That every observation is, in some sense, an interpretation in terms of some theory lurking in the background, has been the focus of protracted discussion in the philosophy of science. Norwood Russell Hanson, the prolific and colorful “Flying Professor” of philosophy, drew attention to the theory-ladenness of observations [52] while Kuhn lent a great deal of weight to this view, within his own terms of discourse, by underlining how the perception of a scientist depends on the world of beliefs and theories she resides in ([76], chapter 10). It is now commonly accepted that observations are indeed conditioned by conceptions of a theoretical nature ([20], [88]) acquired in past experience, and that there is no great distinction between the observable and unobservable parts of reality.

However, the theory-laden nature of all observations has lent itself to other interpretations as well, notably an interpretation espoused by anti-realism in general, and relativism [14] in particular. If observation is theory-laden, then there is no independent determinant of reality since it is through observations that we are supposed to come to our understanding of reality. In particular, various alternative conceptions of reality are possible, depending on the theoretical framework we choose to employ

which, in essence, is the point of view of relativism. Of course, almost everything can be given an interpretation in accordance with our chosen point of view, and the idea of theory-ladenness is no exception. For instance, one can say that our conception of reality is constituted by a complex interplay of facts of observation, originating in an independently existing reality, and the framework of prior beliefs and conceptions that we constantly make use of in interpreting the observations, which would be closer to how I should like to express things.

And, as for relativism, I should not be disturbed by the fact that there can be various different interpretations of any given part of reality, corresponding to various different ways we make sense of facts of observation in terms of our prior beliefs and concepts, so long as we recognize that these interpretations are all oriented towards the *same* reality — the latter continues to exist and to evolve without regard to how we interpret it. Moreover, the framework of beliefs and concepts that we use in interpreting facts of observation cannot be just anything, constructed extempore at the bidding of our whim, because any such framework is built in a protracted course of experience involving innumerable acts of inference and interpretation in the past — ones moreover, that have been tested against earlier facts of observation. It may very well be that there are more than one such frameworks, depending on culturally acquired differences in modes of thought, but that still does not make these arbitrary, with little commitment to explain an independently existing reality and, in particular, to explain further facts of observation as these are found in the course of time. It is here that the various different theoretical frameworks will have to face a reality check when some of these will prove ineffective, some less so and, perchance, one among these alternative theories will get transformed into a broader and more powerful theory to take on the continuing challenge of reality.

But we should not digress too much. We adopt the position that, in a certain sense at least, the table and the electron are equally real and, on the other hand, our perception of either is akin to a theory. But, what does this actually mean? Does this make them amenable to unambiguous description by our theories? Here lies an open terrain that is not so easy to map. Because here, precisely, lies that gap between reality and our

conception of reality. There must be something out there that is causally linked to my perception of the table or to my conception of the electron. But that something is — I am rather tempted to say — quite fathomless. It gets only very partially and incompletely registered or mapped into my perception or conception or theory, whatever you will. Consider, for instance, the table first. Is it that solid object made of wood that has four legs and a plane top that stands before me? Or, is it that collection of atoms or molecules that do not have precisely defined positions and are incessantly in vibrational motion around their mean positions, held together by invisible bonds and separated by vacuous spaces? Or is it a still more nebulous thing comprised of protons, neutrons and electrons, with all these corpuscles engaged in a crazy whirlwind of a dancing exercise? Or, even focusing on a single electron, what can one make of this astounding speck of reality? Is it a particle with a certain absurdly small mass, a certain quantum of charge, a certain 'spin' (in the jargon of physics), and so on? Or, is it, once again, something completely different, like, say, a certain state of a wave field?

Observations are partial: multiple layers of description

Whatever reality is out there, it seems to be really pretty inscrutable! Every single bit of reality has multiple levels of description. And each of these multiple levels invokes a picture of that bit of reality utterly different from the one evoked by another level. You will find realists (many of them scientists, really) who refuse to be intimidated by these multiple layers of reality and rather take it in their stride without making much mystery about it, telling you that this is the most natural thing to be expected of science which goes ever deeper into the description and explanation of the real world, discovering ever more fundamental modes of description. Indeed, probing nature at greater and greater depths is the specific business of the scientists (we are principally talking of physics and physicists here, but our discourse, generally speaking, will be inclusive of other disciplines as well), which is why they are quite so nonchalant about these multiple levels of description.

Here, though, is a question that requires serious thought. Our conception of the world

goes on changing like, for instance, from the conception of the table as a single object to the table as a collection of atoms and molecules, to the same table as a collection of protons, neutrons, and electrons, and so on. Evidently, none of these conceptual changes was accompanied by a corresponding change in the world itself, which means that the thing out there that we refer to as ‘the table’ remains its old self, whatever that self is. While the world does not change, our conception of it changes, and changes rather radically. This, of course, testifies to the existence of the gap between the world and our conception of it. But, granted that the gap is there, what can one say of the successive stages that our conception of the world passes through?

Even as I have depicted the successive stages of our conception of the table as differing drastically from one another, the actual process of change in our conceptual world is, in most cases, more gradual. Conceptions don’t change overnight. There was a more or less prolonged stage when people did not know of electrons and protons, but had a fairly good idea that matter is made up of atoms and molecules. The concept of electrons, protons and neutrons, and their role in the structure of atoms and molecules took a long time developing and maturing. But at the end, people did arrive at a theory of structure of matter that looked radically different from the theory based on atoms and molecules as the ultimate building blocks.

Here, indeed, is a tricky question. Science is supposed to give a true picture of the world. Which of the successive conceptual stages I mentioned above, are true, in the sense of being a correct representation of the object out there we have been referring to as ‘the table’? Evidently, not all of these conceptual representations can be true at the same time. Okay, so can we say that none of these is true by itself in any absolute sense, but that these are more and more accurate representations of ‘the’ truth? This is the view that most realists, and most scientists, appear to subscribe to.

We will have more to say on this later in chapter 9, where we will see how complex the concept of *truth* about nature is. Briefly stated, truth has two aspects to it, one relating to the mental process of arriving at truth — a process that results in an interpretation

of some part of nature, and not in an exact description of it, and the other relating to nature itself, to which the truth pertains. The first of these two aspects can, in a sense, be said to belong to the domain of epistemology, and the second to that of ontology.

I have referred to the irreducible gap between nature and our conception of it. It is precisely because of that gap that our scientific theories, however successful and accurate these are, are interpretations, and not exact descriptions of nature. The gap relates to the fact that there is no mechanism by which nature as a whole imprints an exact copy of itself on our minds. Instead, numerous signals originating in the world out there, transmitted through various channels of observation, stimulate our senses and these stimulations then interact with conceptions stored in our minds that have been produced in a long series of past experiences. It is due to this specific manner of processing of incoming information in the light of past experience, where the incoming information is filtered in a selective manner, depending on current goals and our cognitive and perceptual capabilities, that all our inferences and theories assume the nature of interpretations. *The contrary and complex nature of truth resides precisely in the fact that the vehicle of truth is an interpretation rather than an exact description.* It is this that lies underneath all the clashes, conflicts, and turmoil in everyday life as also in science, where men have to exercise their judgment to approach truth, but have no magic wand of judgment at their disposal.

In scientific investigations, however, magic wand of sorts is, in a sense, provided by nature itself because here a standard of judgment emerges through repeated processes of experimentation, observation, and confirmation. While none of these processes elevates a scientific theory from the status of an interpretation to that of objective truth in any absolute and logical sense, we will assume that the scientific 'process', to all intents and purposes, is capable of interpreting parts of nature within given domains of investigation in an impersonal way, i.e., in a manner independent of the process of interpretation occurring in the minds of individuals and groups of individuals. However, so far as an understanding of nature as a whole is concerned, a scientific theory, in spite of the elaborate process of justification that elevates it to the status of an impersonal interpretation, remains an interpretation nevertheless, since it captures

only some limited aspect of nature, and is contextual, where the idea of context is again a complex one, involving 'internal' and 'external' aspects, as we will see below. As a consequence, the evolution of scientific theories is not a smooth one since there occurs a dramatic restructuring of a theory as the external context of observation and interaction with nature gets changed, and also as the internal context, made up of the existing framework of our ideas, evolves with time. In other words, the way a succession of scientific theories captures the truth about some part of nature cannot, generally speaking, be described in simple terms as a cumulative approach where some ultimate truth is approximated more and more closely.

The existence of objects, and their relatedness to one another

In speaking of the ontological aspect of truth, it is really not enough to speak just of the existence of the world, along with all its parts that we observe either directly or indirectly. Mere existence does not tell us anything. What is of vital importance is the relatedness of objects in the world to one another, because it is the relatedness that tells us how objects behave, what rules, if any, they conform to, and how we can act back on parts of the world. It is the relatedness of an object with other objects that tells us of its qualities or properties, and the properties are precisely the aspects of an object that give it distinction and identity. The properties of the table identify it distinctly from the chairs arranged around it or from the floor on which it stands. So, the correct ontological statement would be that the table, the chairs, the floor, and all such things exist and are related to one another in a manner independent of our mind and of how our mind attempts to know of their existence and relatedness.

However, the existence of objects and their relatedness to one another, make up a single whole. It is only provisionally that one can distinguish between the aspect of existence of natural objects and that of their manifold correlations, appearing as their properties and mutual interactions, of potentially infinite complexity. The distinction between the existence of things and their properties and mutual interactions arises

from our experience of nature, and is also routinely accepted in scientific practice where an entity with some definite properties is first postulated to exist, which is then assumed to interact with other entities similarly postulated. However, the properties postulated to characterize an object are themselves indicative of some of its previously acknowledged characteristics — ones that have already been learned about it.

Theories as condensed descriptions of inter-relations

All our inferences and theories are actually designed to tell us how the things in the world are related to one another. In the context of our attempt to know of this world, it is the relatedness that is of primary importance. Another way of saying this would be that we come to know of the existence of objects in terms of their qualities and properties. And, in order to learn what the properties are, how an object is related to other objects, we ourselves interact with it, either directly through our senses, or by examining it by means of other objects, by examining its relatedness to these other objects. In any case, the properties revealed to us this way depend on how we set up the interaction with it, the context in which the interaction is set up.

To those of you who have an acquaintance with quantum theory, the above paragraph might look like a reference to the notorious measurement problem, because the measurement problem is crucially dependent on the context defining the measuring process. However, I do not specifically refer to the quantum measurement problem in our discourse here. The type of context-dependence I refer to is relevant to the observation of objects in general, where the objects are found to reveal aspects of nature in a manner that depends on the context of observation. The context-dependence of observations in quantum measurements constitutes a very distinctive instance where microscopic systems are found to differ radically from large scale (macroscopic) ones. It seems to me to be a plausible position to adopt that quantum contextuality — I use the term 'contextuality' to denote 'context-dependence') is something that is continuous with, and at the same time constitutes a very special case of, the contextuality inherent in the observation of natural objects and processes in general. But we will

not confine ourselves to stray remarks, and will initiate our discourse on contextuality in the next section, postponing further discussions to later occasions.

All observations and theories are contextual

What I mean by contextuality in the observation of natural objects and processes can be explained with the help of a number of examples. The context in the observation of an electron in an ionization chamber when the electron possesses a relatively low energy differs from the context where the observation is made in the setting of a scattering experiment or in some other similar setting where the particles involved in the experiment possess a sufficiently high energy. One then comes to conclude that the electron is not even a well-defined particle (or a particle-wave in the quantum mechanical sense), but is a certain state of a field. Low energy and high energy contexts are seen here to reveal radically different aspects of nature. It does not really matter that this example relates to the quantum world because it makes us aware of the general fact that any observation of an object or an entity gives us a conception of just a cross-section or a tiny piece of nature, a cross-section determined by the context of the observation. What is objective is the cross-section in question, and the concepts and theories built up for the purpose of representing that cross-section can be faithful and accurate, to whatever degree possible, only within the context — the ‘external’ aspect of it.

On the other hand, at any given stage of development of scientific theories, there exists a vast web of ideas and concepts related to one another, where these concepts are clustered into more or less coherent wholes — the theories, that give us a picture of nature, a map to tell us what things there are and how they behave. However, this vast and complex web is a dynamic one. Concepts and theories — resources that constitute the ‘internal’ context of observation and theory building — get transformed depending on the way science looks at nature. In other words, there occurs an evolution of the external and internal contexts caused by a broadening of the scope of observation of natural things and their behavior, as also by a restructuring of the conceptual web. New aspects — or ‘dimensions’ — of nature are thereby captured in the concepts and theories,

amounting to a radical transformation in the picture of nature that the sciences build up for us.

A good way to illustrate the idea of the context is to refer to the practice of engineering drawing where engineers and architects prepare plans, elevations, and sections to represent aspects of a complex three dimensional object, say, a grand architectural structure, on a two-dimensional drawing sheet. These appear as distinct drawings on the sheet, where each of these drawings represents (possibly with quite exquisite faithfulness and accuracy), some aspect of the structure in some context of representation. Thus, there can be several elevations for observations from several different sides (each, therefore, constituting some particular context of observation), while a top view (a different context altogether) gives the plan. And, additionally, numerous sections may be necessary for a detailed description of what the interior looks like and what the internal structure of each interior object is (like, for instance, whether a pillar is hollow or a solid one). When done by architectural experts on the basis of detailed survey, each of the drawings correspond to some aspect of the structure to be described faithfully and accurately (hence, objectively, though with some unavoidable error which we need not take cognizance of for the present), but none is a representation of the object itself. The object itself, which is most definitely 'out there', appears in our concept only as fractional representations which we continually attempt to synthesize in our mind, and the real question that then comes up is, in what relation does the synthesized concept stand to the 'reality' since, in a very definite sense, the former is an interpretation of the latter.

While the program of science is similar to the architects attempts at drawing up plans, elevations, and sections, it differs in one fundamental and profound respect: an architectural structure is an object of which the description by means of plans and elevations can, in a sense, approach more and more closely the 'object itself' (though, even this is arguable: but we will not be too insistent on finer issues) since the structure has (again, arguably) only a finite number of distinct aspects, or 'dimensions', to it. In contrast, nature is inexhaustible in its aspects ('dimensions') and presents an unending succession of facets for scientific theories to capture and to describe.

I will now pause to draw a few significant conclusions from the example of the electron and of the architectural marvel. The architecture is out there for all to see and to wonder at, while the electron (or whatever it is that we call an electron) is hidden from our bare senses, but there are common aspects to our conception of the two. First of all, the architecture is not really for all to see and visualize since there are hidden aspects in it too — and I don't refer here to the fact that it is made up of electrons and protons or of quarks or whatever, but to the more mundane fact that some of the pillars may be internally hollow while some others are solid, some of the stones are fitted with gems, and the surfaces of some of the domes are serrated, and so on. Every time some new aspects of the structure are revealed to us in some new context of observation, a new synthesis is made in our mind that alters and replaces the conception that was previously there. What is important to note is that this alteration is a qualitative one, and not just an improvement in accuracy and faithfulness and, in this sense, is a replacement. Contrast this with the other type of improvement of representation that an engineer or architect often effects, namely, an improvement in the accuracy of a drawing within any given context like, say, altering slightly some particular angle in a plan or elevation drawing, which also constitutes a replacement, but of a relatively simple nature, namely, a quantitative one.

Another illuminating analogy (or, an instance, if you will) as to how our conception of nature gets conditioned by the context, is obtained by referring to the way we appraise the personality structure of an individual. Our 'reading' of her depends on the various different circumstances in which we observe her response to environmental inputs and, additionally, what our current theory of mind is — in particular, what our current appraisal of her is, because it is the current appraisal that acts as a conditioning factor to how we set up our interaction with her. There are infinitely many aspects to her personality that even a lifetime of probing and interaction will not bring out to the full. Indeed, the very concept of a personality to discover and describe, is not a substantive one. We discover a person bit by bit, but it never amounts to discovering the person as such, whatever the 'person' stands for. A routinely known aspect of her response to her environment may undergo a spectacular change as some components of that

environment get changed. What is more, our conception of the person also undergoes sea changes as we get to know her across a spectrum of circumstances.

Thus, the epistemological question of our conception of the existence of things is related in a complex manner to the ontological question of their existence in reality. My conception of the existence of the table or of the electron is dependent on the context of my observation and experimentation, where the context reveals only a slice of the reality I am looking at, made up of only a few of the infinite number of aspects or 'dimensions' of the latter. And that conception gets altered quite spectacularly in successive stages as the context gets broadened more and more, with an ever increasing number of 'dimensions' or 'layers' of the entities brought under the scanner. Finally, this external aspect of the context is supplemented by and interacts with the internal aspect I mentioned earlier, the two together making up the overall context in which we make observations and build theories in our engagement with nature.

Theoretical concepts: greater and greater depths in a layered description

What is more, the gap between the reality out there and our conception of that reality gets even more intriguing when one considers objects and entities, not only in their aspect of existence, but of their quality, correlation, and interaction, of the mechanisms underlying their multifarious interactions. As I have mentioned, the question of the existence of a table or of an electron is inextricably woven into the question of how it relates to the rest of reality, what the properties of the table or of the electron are, what the mechanisms of its interactions with other entities are. The world exists as a whole and it is only in our minds that we perceive, first, objects in individual existence and, then, the behavior of objects in interaction and evolution, manifest in their innumerable qualities.

For instance, consider a pair of particles A and B. Now, the mere existence of A and B is a rather trivial matter to worry about. One has to look at how these two are related

between themselves, how they interact, how they influence one another, and it is this that makes our conception of their individual existence really meaningful. And it is here that our concepts take a spectacular dive into depths that may in the end prove to be fathomless. Because, now we are concerned not just about the existence of the two particles, but about how best to describe their relatedness. And, just as the existence question is no less conceptual than observational, the relatedness question is equally, if not more, so.

Continuing to refer to the two particles A and B, each of the two can be said to move through a succession of positions depending on their initial positions and velocities, and on their mutual interaction. Physicists describe the latter in terms of a certain interaction potential, which determine the trajectory of each of the two, a curve in three dimensional space. Now, these are concepts at a different level compared to just the concept of either particle existing by itself or even to the concept of the joint existence of the two particles. Concepts of interaction and trajectory may be, in this sense, said to be derived ones. And conceptual complexities continue to crop up at every small step that one takes in describing the behavior of the two particles in their aspect of relatedness, where concepts derived in successive stages continue to make their appearance. Thus, for instance, it is not enough to speak of the trajectory as a succession of positions occupied by either particle at successive instants of time because, of equal importance is the sequence of velocities, or their momenta, at the successive positions. One thereby gets to the idea of a trajectory or a curve, not in our familiar three-dimensional space, but in a six-dimensional one (the so-called phase space) for either particle. The dizzying journey does not stop here, for, one has to refer to a twelve dimensional phase space so as to properly describe, in general terms, the joint motion of the two particles under their mutual interaction. But, hold on. Perhaps we are getting carried away in our newfound enthusiasm with ideas. Why should one go over to a six dimensional or a twelve-dimensional space, and why should it not be possible to keep things concrete and describe the motion of two or, if we like, any number of particles, in our good old world of three dimensions?

It is certainly possible to describe the motion of any number of particles by referring to

just the three dimensional space and a set of mathematical functions. Indeed, the higher dimensional spaces are convenient geometrical constructs for describing features of the functions that one needs in the three dimensional description, but before dismissing these in favor of the mathematical functions in the three-dimensional setting, one has to recognize that the mathematical functions themselves are constructs of a similar nature. However, the higher dimensional spaces force themselves more and more into our reckoning as we go on to describe more and more complex aspects of the mutual influence of the two particles on the motion of each other, or to consider the interactions of a larger number of particles with one another. In attempting to describe the motion of a system of particles, there arise questions of greater and greater depth — ones of greater and greater intrinsic complexity — questions that are virtually impossible to settle in the setting of the familiar three dimensional space. Phase spaces of arbitrary numbers of dimensions then become a necessity, and instead of the interaction potential, one then more conveniently makes use of the *Hamiltonian function* of the system of particles under consideration.

The big question: how do theories correspond to reality?

It will not do to go on here with examples of how abstract and labyrinthine our theories of the mutual interactions of objects and entities can be and how much of a success those theories can prove to be. What is of importance to note here is that, the vast and intricate web of interrelations between objects and entities in nature is necessarily described in terms of theories of a more or less complex nature, theories involving concepts and constructs that appear to be remote from the things existing out there. It is, of course, meaningless to ask whether a Hamiltonian function exists in nature. And equally meaningless to ask if there exist twelve dimensional phase spaces or ones of even higher dimensions. Our description and explanation of interrelations among objects is utterly and incorrigibly conceptual and theoretical, involving theoretical constructs. Granted that such constructs turn out to be immensely successful in answering subtle questions and explaining deep and complex phenomena, are these constructs anything more than mere instruments that help us set up explanations, without hav-

ing something in them that can be identified as truth ? Do the theoretical constructs correspond to some mechanisms intrinsic to nature? Do the sciences converge in ever greater measure to these intrinsic mechanisms that are 'out there' for us to discover?

I am not sure if anybody can produce a definitive answer to this. Evidently, the theoretical constructs are not mere instruments assembled fortuitously, helping us with explanations of phenomena so that we can, by making use of these constructs, identify and predict regularities and correlations with uncanny success. That would be miracle indeed. Like many other realists, I don't opt for miracles. But then, what is there in the hidden mechanisms of nature that these complex theoretical constructs represent, or correspond to? A theory, along with its constructs, must in some way hit upon 'correctly' some mechanism inherent in nature so as to be overwhelmingly successful in some area of scientific inquiry. As I see it, this is the central question that scientific realism has to figure out with some clarity before one can adequately demarcate between the points of view of realism and anti-realism.

The viewpoint of scientific realism has many facets, of which this question relates to an important one. In the end, however aligning oneself with realism or anti-realism is not a matter of hard logic, but one of choosing a position on the basis of shared perceptions and insights. In real life, we often adopt decisions that cannot be proved to be correct on logical grounds (this, indeed, is the hallmark of inductive inference), but we can still try to evaluate those by referring to various circumstantial and contextual factors. Deliberations and judgments are no less important in this world than logical proofs.

The No-Miracle Argument was made famous by the American mathematician-philosopher Hilary Putnam who presented it in support of the point of view of scientific realism in a paper on the philosophy of mathematics ([101], a collection of essays). Bas van Fraassen, who has countered the point of view of scientific realism (or, more precisely, of the point of view that recognizes successful scientific theories as being endowed with truth) with his own constructive empiricism, has dubbed it as the 'ultimate argument' for realism [123], perhaps implying that, tragically, there is no better argument than this in support of realism which, strictly speaking, is no argument at all. Fraassen has advanced his own explanation of the success of scientific theories one where success

is not much more than the outcome of a process of 'survival' by competition: a theory survives the competition from other theories because it happens to have "latched on to actual regularities in nature". Notably, he accepts here the position of metaphysical realism, which speaks of a reality to which all our theories are addressed

I should mention here that we are now referring principally to mathematical and, in a broader sense, theoretical constructs in science, that are mostly encountered in theories in the physical sciences. The uncanny effectiveness of these constructs of exquisite mathematical depth is indeed a matter of genuine puzzlement when one pauses to think over these. The Nobel winning mathematical physicist Eugene Wigner [127] famously expressed this in an article entitled 'The Unreasonable Effectiveness of Mathematics in the Natural Sciences', and several other leading mathematical physicists have expressed a similar sense of wonder and mystery in the fact that mathematical structures, predominantly in the nature of mental constructs, are found to be relevant in the explanation of mind-independent natural phenomena. However, the issue is of more general relevance in the philosophy of science where theoretical constructs are effectively made use of in explaining observed aspects of reality in all the scientific disciplines.

What is more important than proclaiming oneself as belonging to this or that camp in philosophy is to try to gain genuine understanding, even without the benefit of logical proof, of how things operate in this world of ours. And the question of how the deeply theoretical constructs of science correspond to, or represent, the mechanisms of nature, stands out as one of paramount importance.

It is not enough to say that the theoretical constructs are 'true' in the sense of corresponding with features and relations existing in nature because, to me at least, this sounds much like a statement of faith. One has to understand with some clarity what this correspondence actually consists of. Can one say, for instance, that there is a correspondence between relations among elements of a theory and the web of relations existing between natural entities? While this is also vague and unsatisfactory, I will use this as the base camp from where one can have a view of the really intriguing questions concerning the world and our existence in it.

A metaphorical description of how mathematical constructs turn out to be relevant in explaining the mechanisms inherent to a mind-independent reality is that the mind, taking in cues from an external reality, spins out webs of its own, based on rules of logic and mathematics, and some of these webs latch on to aspects of reality. Of course, the webs spun out with threads of logic and mathematics, all hang together, and so do all the various parts of reality and all the mechanisms inherent in these.

Inferences and theories are produced in a psychological process in which beliefs of various degrees of generality play a significant role, as we will see later in this book, but as they acquire a finished form, they pass through a process of justification, both in the mind of the individual and in the collective mind of the scientific community, though 'justification' and 'discovery' are often inextricably mixed with each other (in this context, see [38], chapter 15). An aspect of justification of overriding importance is consistency — the quality that makes our theories and concepts hang together. Our beliefs, however need not be consistent — some beliefs resist justification and consistency check. Remarkably, the belief system of an individual or of a community still hangs together. This is because *beliefs are tied together with emotions* — this we will have a look at in chapter 6 below.

Summary: the issues of relevance

Before I go on, I will summarize for you what we have had so far the position I want to adopt as also the issues that it leads up to. For the sake of easy reference, I will make up a list.

1. There exists a mind-independent reality: this is a matter of inference from our experience, and not one of logical certainty; there exist alternative points of view, any of which one may adopt, depending on how one summarizes and interprets ones experience in life and the experience gained by mankind in the course of history.

Our existence in a real world entails innumerable problems and conflicts

that we have to cope with and move through; the responsibility lies on us to understand and explain the multitude of events into which we are thrust, including social interactions with people around us. For this, we make use of our naturally evolved inquisitiveness and capacities of inference. We interact with the reality around us and form hypotheses and theories in order to explain why things happen. This is vital for our continued existence, and continued autonomy and authenticity in that existence.

2. In explaining the inner mechanisms of nature by means of scientific theories, we arrive at concepts that involve entities that are not directly observable — ones whose existence we infer by indirect observations and reasoning. However, the existence of the unobservable entities is not to be doubted solely because of this. In the course of development of scientific theories, the inferences about these entities are made more and more reliable, and little doubt now remains regarding remotely observed entities, about the existence of which strong doubts were entertained in the past.

At the same time, inferences may be fallible, and it cannot be ruled out with absolute certainty that some of the unobservable entities assumed to exist in our current scientific theories will have to be defined in a new light in days to come. Indeed, the existence of a certain entity, such as the electron, is quite distinct from our *description* of it, and the latter can and does change radically from time to time.

3. While any object of nature exists and interacts with other objects independently of our conception of it, that conception keeps on changing dramatically, as the context in which we interact with and observe nature gains in depth and breadth, and as our conceptual framework itself gets enriched. Our description of things and processes in nature, and our explanations of phenomena, is effected by means of theories, where the latter are generated in a cognitive process. Broadly speaking, that process originates with signals and stimuli from objects in interaction, and then proceeds through a num-

ber of complex stages: the signals and stimuli generate internal responses (neuropsychological, emotional, somatic) and these responses initiate further processing where our past experience in the form of beliefs and concepts assume relevance. The consequence of all this is that our theories describe parts of nature in a selective manner, in the form of interpretations.

4. All observations of facts of the world are fundamentally cognitive in nature, and the cognitive processing of one and the same fact of the world may result in various different interpretations of it. Scientific theories are not determined solely by facts of nature — the latter only act as the causal origin of the former. The theories are formed in a cognitive process in the form of a selective and purposive interpretation, being conditioned by past experience, and by past beliefs, conceptions, and theories. Observations are theory-laden.

5. Scientific theories constitute a continuation of inferences we make in the course of our daily life, where these inferences provide the basis for our decisions and actions. In the making of these inferences, we interpret the world around us by making use of our cognitive abilities and, at the same time, keep on enriching and transforming that interpretation as the context of our interaction with the world changes and, at the same time, as our overall conceptual framework gets enriched. What is special in the case of scientific investigations is the great emphasis on repeated cross-verification against facts of observation, on rigorous consistency checks against a vast web of currently existing concepts, and on discourse among members of a community of fellow scientists where rules of inference agreed upon by the entire community are made use of. This tends to make scientific theories free, to all intents and purposes, of relativism, in contrast to inferences of a non-scientific nature, where interpretations often vary from person to person, and from one group of persons to another.

This aspect of enhanced reality check notwithstanding, scientific theories con-

tinue to be in the nature of interpretations of reality, since these capture only certain aspects of nature through a cognitive process as outlined above.

Even though impersonal to a large extent, scientific theories bear the stamp of their origin in the cognitive processes of men and of groups of men. The course of development of the theories is conditioned by modes of thought of communities and by cultural traits and latent beliefs of societies.

6. Scientific theories aim at providing us with effective descriptions of entities and mechanisms constituting reality, an effectiveness that is supposed to imply truth. However, the concept of truth is a complex one since, on the one hand, it relates to a mind-independent reality and, on the other, resides in inferences and theories that act as vehicles of it, the latter being produced by cognitive processes in the minds of men. This contrariness results in a fundamental tension that cannot be resolved by logical discourse alone. However, that is how everything in life turns out to be — no concept is ever pure and free of contrary aspects, least of all the concept of truth.

This, however, does not make invalid the concept itself. Truth is not delivered to us ready and tied in a neat package. Mankind has to struggle for it, and struggle with all intensity and commitment to achieve authenticity, as it has had to do in the past.

There is no denying the irreducible gap between Nature and its workings, and the model world that science assembles for us in the form of the great web of concepts and theories, that makes realism a matter of a *stand* or a *viewpoint* that we adopt in our own life process — a viewpoint arrived at in experience, and not as a logically compelling conclusion. At the same time, the gap relates solely to the fact that our perception of reality is cognitive in nature. *Nature does not know of any such gap*. Signals originating in the world interact with our cognitive apparatus, and we as cognitive agents, on reflecting upon our cognitive process, become aware of what we describe as a gap. Looked at from outside the cognitive process (in so far as such a thing is possible), the whole affair

is nothing but an interaction between parts of the world and our cognitive apparatus (with all its ‘software’ of concepts and theories) — just another instance of the vast web of interactions going on between parts of nature. The mysteries of epistemology arise only as we make an abstraction, for our own benefit, of questions of epistemology from ontology.

What is more, the cognitive process itself knows of no great distinction between the pre-scientific and the scientific. Our web of scientific theories has parts of it built up through experiences gained in non-scientific and pre-scientific practice, and the rest is assembled by scientific exploration, where inter-personal differences in interpretation are sought to be minimized. This lends a new complexion to scientific theories that can be likened to an edifice built upon a substratum of loose ground, not as coherent as the edifice itself. But that is precisely how all edifices are built.

“Science does not rest upon solid bedrock. The bold structure of its theories rises, as it were, above a swamp. It is like a building erected on piles. The piles are driven down from above into the swamp, but not down to any natural or “given” base; and if we stop driving the piles deeper, it is not because we have reached firm ground. We simply stop when we are satisfied that the piles are firm enough to carry the structure, at least for the time being.” (Karl R. Popper, in *The Logic of Scientific Discovery*, quoted in [82]).

We will now have a look at all these issues relating to the scientific process and its relation to cognitive activities of men, but there is one other thing to take note of before we can proceed.

The questions that are coming up are not as much in the domain of science as such, as in that of philosophy of science. If the theories of science constitute a distilled essence of our experience with reality, then philosophy of science, in turn, constitutes the distilled essence of those very theories. If the scientific concepts and theories are arrived at by acts of inductive inference, then the points of view one adopts in philosophy of science are, in turn, arrived at by induction of a higher order.

Philosophies are produced not by the passive application of universally accepted rules but by a choice from among alternatives, where a choice cannot be right or wrong but can only be more plausible or less so. The alternatives are, fundamentally speaking, not many, because philosophical ideas are made up of only few building blocks when compared with the vast and awe-inspiring array of concepts the scientist works with. But the scientist has an arbiter standing in front of her to whom she and her fellow travelers defer in judging the worth of their theories — Nature herself. To the philosopher, on the other hand, Nature does not proffer any yardstick for judging right or wrong — she only smiles enigmatically (however, even for the scientist, the yardstick does not come ready-made; but more of that later). Science creates a new world for itself every time the horizon of mankind's interaction with nature expands; philosophy only recreates the world — a world built in contemplation.

What makes philosophy a really challenging exercise is the innumerable nuances that the few basic ideas are combined with. Where the scientist works with a great variety of building blocks, erecting structures of awe-inspiring complexity, the philosopher works with clay — only a few colors of it — to which she gives innumerable shapes. The philosopher of today works with the same old clay and the same old colors that the philosophers of antiquity worked with, but the shapes made up by the two differ much, because the edifice that science assembles goes on being built, acquiring new additions of ever-renewed shapes. While the edifice of scientific theories is built in successive episodes upon the existing edifice and takes on ever-expanding fantastic forms, philosophies are re-built. There appear cycles in philosophy, where old ideas are worked anew, worked in new contexts, because the philosophy of science works on scientific theories — theories that keep evolving in fantastic ways. And, precisely because of this, philosophical practice is forever filled up with polemic, with exchanges between competing camps. In contrast, scientific practice involves polemic only during the phase of emergence of a new theory, or at a time the existing theory proves sterile in solving new problems while no definitive shape of a successor has emerged. Once a new theory proves its worth, polemic is replaced with feverish and predominantly cooperative work — perhaps competitive as well — in the journey forward. In philosophy, you can say nothing really new, you can only adopt a different way of looking at

things. And, it is precisely because of this that there cannot be sharp and lasting demarcation lines between the various camps that philosophical thought is divided into. Old demarcation lines dissolve and new ones appear. Contrasting ideas appear to be irreconcilable and mutually exclusive only when these are allowed to be bled of life by being defined with empty words — words that don't really apply to the context at hand.

While one is to choose between realism and anti-realism, it is by no means a matter of a simple dichotomous choice between two mutually exclusive clusters of ideas. There do exist distinct cores to these two positions in philosophy, but the cores do not exhaust the entire philosophical terrain and it really makes no sense to carve up the entire terrain into mutually exclusive clusters of ideas. On the contrary, it may be immensely fruitful to think of novel combinations of elements of ideas that have so long been assumed to belong to one or the other of these mutually exclusive clusters. Indeed such novel blending of ideas picked out from what previously appeared to belong to irreconcilable camps, may be more in consonance with the developments eternally taking place in the sciences where, again, old frontiers dissolve and new areas of exploration emerge at an astounding pace. Indeed, explorations for a workable position in the philosophy of science are to be conducted in close affinity to developments in the sciences, taking care that such explorations are constrained by these developments, by the successes and failures of scientific theories. This, broadly speaking, is the naturalist position in the philosophy of science.

This has been the way Philip Kitcher's position in the philosophy of science appears to have evolved over the years. For instance, [33] indicates how Kitcher's realism has undergone a process of moderation as his views have broadened during the years between the writings of [71] and [72]. Even at the time of [71], Kitcher speaks of a "vast middle ground" between extremes in issues relating to scientific realism.

I will, then, close this chapter on the note that we will, with these few strands of ideas introduced here, again face the question we started with: in what sense does science

describe the workings of nature ? As I understand, there is no clear-cut answer to the question, which is to be approached, not head-on, but only from the standpoint of our own cognitive processes, of how we perceive the world. In the end, however, all these make sense only when looked in relation to the question as to *how we perceive ourselves*. This is the question of values, of our goals and purposes: the ultimate testing ground of mankind.

Chapter 3

The emergence of theories: how are theories constructed and accepted?

The explorers in science: individuals and scientific communities

It seems reasonable to describe this fusion of the personal and the objective as Personal Knowledge.

Michael Polanyi

This book aims at addressing two complementary aspects of science from two distinct, though related, perspectives. The first of the two starts by considering science to be a product of humankind without overt concern as to how the product came to be. In reality, science is done by human beings in their individual capacity and as members of groups of people — each as a member of her own immediate group or community of fellow scientists, as a member of a bigger scientific community, and as one belonging to groups of fellow human beings sharing various different identities and cultural values;

and this constitutes the second of the two aspects of science mentioned above. The first of the two perspectives was adopted in the last chapter. We will adopt the second perspective now. However, we will see by and by that the two perspectives are inextricably woven into each other.

The essential tension: the objective and the subjective

Theories in science are constructs relating to aspects of reality. And herein lies the great tension in the world of science: on the one hand, theories have a causal link to slices of reality, depending on the context of our confrontation with these, and, on the other, these are constructs — ones produced in the minds of individuals and groups. On the one hand, science is to be loyal to nature while, on the other, it is an interpretation of nature produced by people whose world of concepts is separated by an irreducible gap from the ‘real world out there’. I have raised, in the last chapter, the question as to how and in what sense science can be said to be loyal to nature. This chapter will deal with the other aspect — that of the way the scientific theories are generated in the conceptual world of individuals and groups or communities. We will then proceed to a more complete and meaningful discourse where these two aspects are merged with each other, making possible a composite picture of science as an evolving map of nature, as revealed in the minds of people.

Science is supposed to be ‘objective’ not only in the sense of being loyal to nature but in that of being loyal to nature alone, being free of the vagaries of the minds of individuals and of communities of individuals. However, like the first of these two suppositions, the second too is burdened with big question marks hanging over it. This is what we will have a look at now.

But even before we proceed with an examination of how science is dependent on the conceptual edifice already there in the minds of people, I must caution you not to read too much into any one of the two perspectives I have mentioned above to the exclusion of the other. In philosophy, as in everyday life, we do have a tendency of reading too much into a single idea and losing foothold. This, for instance, is what I

CHAPTER 3. THE EMERGENCE OF THEORIES: HOW ARE THEORIES CONSTRUCTED AND ACCEPTED?

refer to by speaking of the big divide between realism and anti-realism. Realists speak of science being, in some sense or other, true to Nature, and are skeptical of allusions to the mind-dependence (and culture-dependence too, but culture-dependence can be looked at within the broader perspective of mind-dependence) of science, while anti-realists have the opposite tendency of emphasizing the mind-dependence at the cost of the allegiance of science to nature. In reality, however, the two are blended into a single whole that we call science. It is only for the sake of convenience of discourse that we can effect a separation between the two, where we allow ourselves to forget the essential fusion of and tension between these.

Science on a pedestal

The first three quarters of the twentieth century saw great and dizzying advances in science where remarkable theories were built and, in a concomitant development, a certain mindset was developed, whereby Science was placed on a pedestal. Science was seen as something transcending human fallibility, as a tightly knit body of knowledge, based on and regulated by the inexorable laws of logic, where the individual human being had no place simply because the individual human mind was not fully 'objective', nor were the judgments and propensities of communities of individuals. There were individuals steeped in superstitions; there were tribes engaged in ritualistic dances and mystical magical practices; there were communities engaged in blind and cut-throat politics; and there were nations at war. Then, there were music, poetry, literature, and the arts oriented toward the innards of the human mind. And, among all this blind passion and subjective explorations, here was Science, the only endeavor truly seeking to reach out to Nature by transcending human subjectivity.

Paralleling all this lofty status granted to science, there were trends in the philosophy of science as well, trends that stressed, on the one hand, the great role of logic in science that made it free of the vagaries of human psychology and, on the other, the epic of Science unraveling the mysteries of nature, of science engaging with reality. To be sure, there were conflicts between dictates of logic and claims of science discovering the real

workings of nature because logic, after all, works in the world of concepts and cannot, by itself, bridge the great divide between that world and the real world, the World of nature. But then, it was hoped, science could make our conception of nature converge more and more to nature itself.

Logic and reality

Philosophy of science in the first three quarters of the twentieth century was dominated by Logical Positivism and Scientific Realism. The logical approach applied the proverbial Occams Razor, with logic ideally playing the role of that razor and, to be consistent in its dissections with the razor, could not, in principle, accede without reservations to the reality of what goes by the name of Reality, especially that part of it which is remote from our senses. Instead, science was seen as attempting to identify regularities of what is sensed and experienced. Scientific realism, in its turn, in acknowledging a mind-independent world, had to grapple with the question of how the mind can arrive at truths about that world — that central question of epistemology. In other words, Logic and Reality resided in distant worlds, and the two could not be united in a consistent manner — consistency being supposed to be what science was all about.

But questions of philosophy are abstruse ones. The general perception of science, in both the common man and the working scientist was that science was utterly logical and that it revealed for mankind the real mechanisms underlying the workings of nature. The two together made up the picture of objectivity of science.

In philosophy as in popular perception, science was seen as something rational because (a) it was logical and thus free of the common fallibilities of the human mind, (b) it was aimed at the truth about nature, and (c) it tended to approach that truth more and more closely.

Logical positivism was viewed variously even by the proponents of the philosophical movement themselves. It was based on the analytic-synthetic distinction originally introduced by Immanuel Kant, where propositions in logic and mathematics were of the

CHAPTER 3. THE EMERGENCE OF THEORIES: HOW ARE THEORIES CONSTRUCTED AND ACCEPTED?

analytic type, while statements with an empirical content were of the synthetic type, which called for a principle of verifiability of such statements. The question of verifiability brings up the issue of confirmation of scientific hypotheses and theories, but no logically compelling method of confirmation could be found. During the later half of the twentieth century, logical positivism gave way to logical empiricism in which the failings of the logical positivist project gradually gained acceptance. At the same time, the overall philosophical framework was sought to be retained, though without any strong unifying thread. Logical empiricism was thus, in a sense, the self-examination of logical positivism. It was this self-examination that made room for the development of the pragmatist and naturalist trends in the philosophy of science, within a broadly realist framework.

An account of the logical positivist and logical empiricist movements in the philosophy of science can be found in [46].

Logical positivism was no longer the dominant current in the philosophy of science by the fourth quarter of the twentieth century, while it was precisely this period when naive realism as a trend was put to question as well. In a sense, the trouble with both approaches relates to that great gulf separating the conceptual world of men from the real world out there, where the two worlds could never be bridged seamlessly. Logical positivism harbored within it the skeptical viewpoint that questioned everything outside the realm of sense data and could not come up with a solution to the question as to how the truth of scientific theories could be established by logic, for which it tried to develop a foolproof theory of induction. If a logic of induction could be developed, then that would justify the confirmation of theories by means of empirical observations. But no such logically sound foundation of induction was forthcoming, and the state of affairs remained pretty much the same as where David Hume had left off, with induction continuing to remain as the ‘scandal of philosophy’ [15].

Realism, on the other hand, grappled vainly with the question of the truth of scientific theories, which it could not resolve in a logically sound way. As in the case of logical positivism, realism got trapped within the cage of consistency — how could the truth of

scientific theories, being objects in the conceptual world, be ever proved to be true of the natural world since the two worlds are fundamentally remote from each other? The last three decades of the twentieth century saw a gradual withering of the lofty status of realism in the philosophy of science.

Michael Polanyi: roots of personal knowledge

The first major current against the impersonal objectivity of science had been initiated by Michael Polanyi, a scientist of great repute who turned to the field of philosophy and put forth ideas of astounding originality, only to be set aside by the philosophical community with what amounted effectively to a condescending dismissal meted out to the amateur.

Michael Polanyi (1891-1976) was a physical chemist of great stature, and turned to philosophy in his later years. The relevance of his contributions to the philosophy of science will, perhaps, take many more years to be fully realized and appreciated. A biography of Polanyi, including his scientific contributions and his economic, political, and philosophical thoughts, written by his long-standing friend Eugene P. Wigner, who was also a collaborator, is to be found in [128] (Wigners account of the scientific contributions of Polanyi is to be complemented by the mention of his path-breaking work on crystal dislocations that opened up the enormously important field of plastic deformations). Wigner and Polanyi had a long correspondence and dialog on questions relating to tacit knowledge, quantum physics, and the mind-body problem, for which, see [67].

Polanyis major work of great relevance in the philosophy of science was put together from a collection of lectures and essays, and was named 'Personal Knowledge: Towards a Post-Critical Philosophy' [94], a major concern of which related to the acquisition of knowledge as the act of individuals and, in particular, to the building of scientific concepts and theories as a personal endeavor, where the major thrust of Polanyis was against the all-consuming critical stance of the logical positivist who was skeptical of

CHAPTER 3. THE EMERGENCE OF THEORIES: HOW ARE THEORIES CONSTRUCTED AND ACCEPTED?

whatever did not conform to the canons of logic. He was one of the early philosophers of science to grasp the significance and relevance of cognitive processes taking place below our level of awareness (“we can know more than we can tell” [95]), or what can be termed pre-conscious and pre-logical stages of cognition. He analyzed the process of formation of conjectures, hypotheses, and scientific theories as these actually arise in the minds of people, in which he can be identified as perhaps the first major naturalist in the modern era. Polanyi, after Freud, was one who delved into the human unconscious as the substratum of conscious activity, though, to be sure, his thrust differed from Freuds in that his major concern was to look for the roots of human cognition, where the latter has a noticeable conscious component as well.

In looking at what he termed *tacit knowledge* (or what I feel should more appropriately be termed tacit *cognition*) Polanyi underlined the role of *beliefs* entrenched in the human mind where these beliefs endow the cognitive endeavor of the individual with deeply personal explorations and commitments leading to the birth of new concepts, hypotheses, and theories. But this concern of Polanyis with unconscious personal cognitive processes did not make him an anti-realist in the sense of identifying scientific theories as solely mind-dependent constructs having no connection with reality. In fact, Polanyi was a realist even as he underlined the deeply personal aspects of theory building in science. His view of the process of emergence of scientific hypotheses was that it was anchored in reality, though not directly so. Instead, the ‘external’ reality was seen as causing the generation of a huge store of unconscious bits of *clues*, the constituents of tacit knowledge. This vast and ever-growing store of tacit clues was seen as providing the real context in which concepts and theories were produced.

“To say that the discovery of objective truth in science consists in the apprehension of a rationality which commands our respect and arouses our contemplative admiration; that such discovery, while using the experience of our senses as clues, transcends this experience by embracing the vision of a reality beyond the impressions of our senses, a vision which speaks for itself in guiding us to an ever deeper understanding of reality — such an account of scientific procedure would be generally shrugged aside as out-

CHAPTER 3. THE EMERGENCE OF THEORIES: HOW ARE THEORIES CONSTRUCTED AND ACCEPTED?

dated Platonism: a piece of mystery-mongering unworthy of an enlightened age. Yet it is precisely on this conception of objectivity that I wish to insist ” ([94], chapter 1).

Thus, it was essentially the supreme naturalist in Polanyi that enabled him to mount a major two-pronged development in the philosophy of science, a development that, on the one hand, was directed against the skepticism of the logical positivist denying everything outside the realm of logic, thereby bringing the personal aspects of scientific theory-building to the fore, and on the other, sought to establish a causal link between the reality of nature and the subjective world of the individual in the form of clues in an unconsciously held storehouse of knowledge. It is essentially a link of the same nature as that involved in the perception of forms and shapes of objects in our *everyday* observations, such as the recognition of a face or as the identification of a white, oval object as an egg. Our mental recognition of either the face or the egg is actually a *construction* in our mind produced by sensory inputs working in the *context of stored clues*. In arriving at such a recognition our mind correctly *latches on to* some aspects of reality (as van Fraassen would say), much like a key fitting a lock — it would, metaphorically speaking, not fit to just any lock, but only to the one whose inner structure fits its own (the keys) construction — a construction that was dictated by clues supplied by the lock itself (the key was *constructed* by an independent hand from these clues, but it would succeed in opening the lock nevertheless).

In stressing the aspect of hypotheses and theories being constructed in the minds of individuals, Polanyi highlighted that these are, truly speaking, *guesses* — guesses that were not determined uniquely by known data, being generated in flights of imagination, and ones that left open the possibility of *alternative* hypotheses. At the same time, the hypotheses are not pure fancies of the mind since these are *constrained* from two sides — on the one hand, the concepts, hypotheses, and theories are constrained by the sense data, the inputs generated by our experience in the real world, and on the other, these are constrained by the *context* of our already acquired cognitive products, the vast storehouse of knowledge, memories, and clues — partly conscious and overwhelmingly

CHAPTER 3. THE EMERGENCE OF THEORIES: HOW ARE THEORIES CONSTRUCTED AND ACCEPTED?

tacit. The treasury of tacit knowledge and cognitive clues, in turn, is a product of past acquisitions from experience where clues were received of the workings of the world out there, and it is thus within a complex and stratified web of existing concepts, clues, hypotheses, and theories — *all* invisibly linked to reality — that new hypotheses and theories are generated. This is the two-fold nature of newly generated hypotheses and theories — being products of imaginative guesswork constrained by present and past experiences of nature, where clues received from nature are made use of. Imagination and reality — these are the two contrary things that encapsulate the essential tension inherent in the entire endeavor of science, the tension resulting from the fathomless gap between the world of nature and our conceptual world.

“We see here [in scientific discovery, akin to guessing the presence of a burglar in the house at night] a consistent effort at guessing — and guessing right. The process starts with the very moment when, certain impressions being felt to be unusual and suggestive, a ‘problem’ is presenting itself to the mind; it continues with the collection of clues with an eye to a definite line of solving the problem; and it culminates in the guess of a definite solution.”, ([96], p 9-10).

But that gap between our mental world and the real world is not something that we need to read too much into because, simply speaking, it is something that we, as limited cognitive agents, perceive as a divide between ourselves and Nature — everything that excludes ourselves. It is only in philosophy that all this smoke is produced where one tries to ‘understand’ this fundamental divide, because all forms of life including human beings happily continue and thrive without being burdened with this awe-inspiring ‘responsibility’ of ‘bridging’ the gap. In the course of existence, an individual living being engages in various life processes on the basis of myriads of clues received from the external reality as also of clues generated internally. The mental activity of man is one instance of this ongoing life process — one of a very special nature though. It is only when the mind focuses on its own activity *vis-a-vis* the external world that the conundrum of an irreducible gap between the ‘knower’ and the ‘known’ is raised. At the same time, it is this gap — the distinction between the real world and our conceptual world

CHAPTER 3. THE EMERGENCE OF THEORIES: HOW ARE THEORIES CONSTRUCTED AND ACCEPTED?

that irrevocably constrains the way we perceive the mind-independent reality that we are immersed in.

And Polanyi — the naturalist that he was — did not take it upon himself to resolve the conundrum, instead choosing to address the question as to how scientific thought is *actually* generated in the minds of men. In this, Polanyi analyzed in great details a number of aspects of the process of formation of new concepts, hypotheses, and theories, much in the manner of an introspective study and analysis, and came out with interesting and important *clues* to this process. The first ‘clue’, of course, was that the process occurred, in the main, within the matrix of tacit knowledge and was essentially in the nature of a guess-work, though one that made use of clues acquired from the reality that the hypothesis or the theory sought to describe and explain. And the clues were immersed within the *belief system* of the scientist, imparting the process with features of a deeply personal nature. A second interesting observation of Polanyi’s was that the process was initiated and sustained by a commitment on the part of the scientist, triggered by a puzzle, a problem, or an anomaly that the existing theoretical framework failed to solve or explain.

A third observation was that the clues were half-baked ideas or hints that were tacitly available to the scientist, having been produced in a long process of experience and prior attempts — failed as well as successful ones — to solve problems of a similar kind. Polanyi highlighted the idea of *heuristics* — hunches and rules of thumb — enunciated by George Polya in the context of problem solving in mathematics and thereby anticipated, along with Polya, a broad and major trend to take shape in subsequent decades in *artificial intelligence* and, more generally, in the field of understanding and explaining inductive inference in the context of human inferential processes.

Yet another feature to emerge in the course of Polanyi’s analysis of the process of the formation of hypotheses and theories was that it was not a process aimed at grasping the *whole* of nature, of nature in its entirety, but one focusing on some aspect of nature, one that relates to the immediate problem or anomaly at hand (“...an aspect of nature seeking realization in our minds”, [96], p 21) while, at the same time, picking up and

CHAPTER 3. THE EMERGENCE OF THEORIES: HOW ARE THEORIES CONSTRUCTED AND ACCEPTED?

pointing at clues to *hidden* aspects. A hypothesis or a theory proves itself superior to an earlier one and supersedes the latter in this ability of providing a glimpse to hitherto unrecognized aspects of nature, thereby finding its place in an unending succession of hypotheses reaching out to an inexhaustible landscape of nature, awaiting exploration. In this optimistic outlook on theories reaching out to nature with its infinite hidden dimensions, Polanyi once again proved himself to be an early exponent of the naturalist trend in the philosophy of science in the recent era, refusing to address lofty philosophical questions relating to the ‘mind attempting to comprehend nature’ — questions that were formal and contemplative ones that required endless reflection, with little possibility of an acceptable solution. And the naturalist in him gave a new and fresh view of realism too in speaking of the infinitude of nature that remains ever unexplored.

“ A true physical theory is, therefore, no mere functional relation between observed facts, but represents an aspect of reality, which may yet manifest inexhaustibly in the future.”, ([97],p 191).

Finally, Polanyi came out with the remarkable observation that a creative act of hypothesis formation on the part of the individual scientist involved a critical stage in which apparently uncorrelated bits of half-formed concepts undergo a spontaneous coalescence into a new coherent form, where the coherence is once again to be interpreted with reference to the previously existing storehouse of tacitly held cognitive elements and to inputs from the external reality.

“We may follow up our parallel between discovery and Gestalt perception by regarding the process of discovery as a spontaneous coalescence of the elements which must combine to its achievement.”, ([96], p 19).

All these aspects of the formation of new concepts and theories in the minds of individuals will be discussed at greater length in subsequent sections in this book. I mention these here in order to indicate how concepts *corresponding to* aspects of reality can be formed in the minds of men where, as Polanyi pointed out, the process is at once

CHAPTER 3. THE EMERGENCE OF THEORIES: HOW ARE THEORIES CONSTRUCTED AND ACCEPTED?

constructive *and* exploratory — one of invention and discovery. It is this interweaving of apparently contrary aspects that gives rise to philosophical conundrum. I say ‘apparently’ here, because the contrariness is a matter of *our* perception — Nature does not have any innate design of baffling us with this hugely tough puzzle of ‘mind versus reality’. We, as thinking agents, are parts of nature and we are ‘forced’ into making maps of aspects of nature in our minds where these maps are *representative* of those aspects of reality and, at the same time, are fundamentally *skewed* ones. The ‘forcing’ is, of course, part compulsion and part adventure — adventure into uncharted realms of reality where the spirit of adventure is, in all likelihood, acquired in an evolutionary process.

It is the personal and constructive aspects of scientific theories that we will be discussing further in the pages of this book, acknowledging at the same time that these aspects are not necessarily antithetical to the *other* aspect of these theories, namely their allegiance to aspects of an *external reality*. In this, scientific theories are, to some extent, comparable to a portrait of an individual created by an artist — the individual is sitting out there as the live model whom the artist represents in his portrait, but what the artist finally produces is not an exact likeness of the external appearance of the model, so much so that the model herself feels anguished at seeing the finished portrait, complaining that she cannot recognize herself in it. While the artist most definitely anchors his work in the presence of the model ‘out there’, yet the portrait turns out to be the artist’s *perception* of her, including ‘dimensions’ of hers perhaps unknown to herself, dimensions relating to her personality, her psyche, and her immediate mood. *Clues* to these other ‘dimensions’ are received by the artist in ways perhaps he himself cannot define, and the greatness of the artist lies in how his tacit and unconscious self makes use of these clues, and combines these with previously stored perceptions in his mind so as to make a coherent whole that, while being a great work of art, is still a skewed representation of the subject sitting expectantly in front of him.

Thomas Kuhn: the paradigm shift

The remarkable insights of Michael Polanyi, however, remained largely unnoticed in the world of philosophy of science. While constituting a telling criticism of the received image of science (and of human thought in general) these did not raise much of a tremor in that world. But Polanyi was followed by *Thomas Kuhn* who did cause a great tremor that extended from the terrain of philosophy of science out towards remote cultural terrains, to merge into broader cultural movements of the last quarter of the twentieth century. Kuhn's thoughts had quite a significant kinship with Polanyi's in that both had a naturalist trend and both focused on the constructive aspect of scientific theories, at the risk of appearing to undermine the received view of the objectivity of science.

Kuhn's naturalism was expressed in his paramount interest in the *history* and *sociology* of science, where he refused to engage with abstract philosophical problems, instead looking at how scientific theories were actually constructed in the historical unfolding of ideas in the context of specific challenges of a social and epistemological nature, and at how scientific concepts and theories are actually handled by communities of scientists in the course of their professional work. He, like Polanyi, was one of the early representatives of the naturalist trend of recent decades without overtly contending or appearing to be so, both primarily engaged in settling a number of issues in the philosophy of science (issues as perceived by them) without burdening themselves with abstract questions of epistemology, ontology, and metaphysics.

However, one cannot glibly dismiss centuries of seemingly abstract philosophical discourse since philosophical abstraction is by nature akin to scientific abstraction. Questions abstracted away from the endless contrariness of real life lead to fruitful analysis but only if the results of that analysis are continually referred back to the muddled grounds of reality. In a sense, the analytical philosophy of the first half of the twentieth century was the necessary precursor to the naturalist-cognitive turn that came about at around the nineteen sixties and seventies. The great merit of the analytical philosophy was that it never desisted from acute self-analysis.

CHAPTER 3. THE EMERGENCE OF THEORIES: HOW ARE THEORIES CONSTRUCTED AND ACCEPTED?

In his first acclaimed work, *The Copernican Revolution* [77], Kuhn presented a study of how a major transformation in the world view of science was brought about under the influence of factors only partly of an epistemic nature, where major *social-cultural* components exert an equally important influence in the complex of pulls and pushes resulting in the transformation. In particular, he identified stages in the process where conceptual changes occurred *without regard to* evidential support, or, in other words, without overt concern for convergence with observation data (Polanyi also drew heavily from an appraisal of the Copernican revolution with his inimitable philosophical insight ([94], [97])). And then he came out with his major and celebrated work, *The Structure of Scientific Revolutions* [76], where he cogently set forth a number of theses relating to the scientific process that put a question mark to the logical-analytic-philosophical view of an impersonal and objective science undertaking the Promethean task of bridging the chasm that separates mind from matter and cumulatively unraveling hidden mysteries of nature.

Kuhn's name has now become synonymous with phrases like 'scientific revolution', 'paradigm', 'normal science' and 'incommensurability'. However, while his work generated a great stir in professional circles of philosophers of science, scientists, and sociologists, and gradually came to create a remarkable impression in the minds of people much beyond the limits of these professional circles — quite in contrast to Polanyi's work — his bold ideas were not received by professionals with open-armed acceptance. What Kuhn had done was to counterpose a *point of view* to another, entrenched, one, in a breathtaking sweep, which is precisely why his work received a *general* acclaim, far beyond the confines of its avowed subject area. But it did not quite match the scalpel of the professional who took to dissecting his views much like the zealous surgeon who dissects in vain to discover where life lies hidden in the patient on his dissection table. To be sure, there were professionals who found Kuhn stimulating but still, they could not quite accommodate Kuhn within a rigorously justified framework that they were accustomed to. But, despite their demands of precision and rigor, their own field would never remain the same when Kuhn was done addressing from the mountain-top. He was speaking, not so much on this or that specific problem in the field of philosophy

CHAPTER 3. THE EMERGENCE OF THEORIES: HOW ARE THEORIES CONSTRUCTED AND ACCEPTED?

of science, but against the mindset of his *time* — the same dead-weight that was felt implicitly by an entire generation. While a work of such sweeping spirit does not appeal to professionals, it does stir up the minds of people who, at least for the time being, are prepared to leave aside their professional armor.

Kuhn, of course, was not alone in addressing the mindset of his time since there were others who gave vent to the gathering tremor in the world of ideas and viewpoints — those in the likes of Polanyi, Hanson, and Feyerabend, each of whom was to leave a lasting impression in the field of philosophy of science; but none of these others caught the *imagination* of their generation, and of generations to succeed, quite as much as Kuhn did. Polanyi, in particular, raised issues of a stupendous magnitude in significance and implication, but was never noticed much within and outside the circle of philosophy of science, because he did not *connect* with people, was felt to be idiosyncratic, and expressed concerns too much at variance with those of his time.

The question as to why Kuhns work drew more attention than Polanyis has been addressed in [119].

On the irrationality of the substratum

What I want to highlight as the seminal contribution of Polanyi and Kuhn in the context of the issue of the objectivity of science is that both underlined the *constructive* aspect of scientific theories — one in the minds of individuals and the other in the perceptions of scientific communities. And in this, they rose to new heights by refusing to relinquish the connection of the theories to an *external* reality, the reality of nature — one acknowledging this connection quite openly and other by implication. What is more, they both trod what appeared to many as the path of *irrationalism* but what in reality was the path of identifying and looking into the *substratum* of the logical-rational mode of thought in science. In this, Polanyi explored an area distinct from Kuhn's — in the mind of the individual as distinct from the pre-logical perception of the scientific community that was given the name of 'paradigm' by Kuhn. It is the substratum that is apparently at variance with the super-stratum of logical-rational and deliberative mode of thought

CHAPTER 3. THE EMERGENCE OF THEORIES: HOW ARE THEORIES CONSTRUCTED AND ACCEPTED?

and it is the discrepancy between the two modes of thought that is commonly perceived as *irrationality intruding into the world of science*.

The process of formation and initial acceptance of new concepts and theories is rooted in the substratum, where the concepts and theories are constructed, but this process of construction is *constrained* by the external reality and is ultimately aimed at it, the logical-rational super-stratum being, in a manner of speaking, the mediator between the two, i.e., between the substratum and the external reality. This, in brief, is the scheme of things that Polanyi and Kuhn outlined in their work. The distinctive features of the areas of discourse chosen by the two determined, in some measure, the roles that these two were to play in subsequent decades. Polanyi, after his time, gained great acceptance in the fields of cognitive science and psychology because he, more than anybody else, re-opened the door to the human unconscious after the latter was banished from respectable science by the ostracization of psychoanalysis and by the rise of behaviorism. And Kuhn, by contrast, brought in a new wind in the fields of historiography and sociology of science that ultimately impacted on the philosophy of science and, then, merged with emerging trends in broader cultural areas.

With this, I am done with laying the groundwork for the remainder of this book. In addressing the problem of objectivity in science I will repeatedly come back to issues raised in the last chapter and the present, but by then you will have known the setting in which one can meet with the ideas I want to share with you — explicit and implicit ones to be found in the literature put together by researchers.

Chapter 4

Inductive inference

Induction: the guessing game of life

Life is an endless process of guessing, and survival and progress means *guessing right*. Guessing essentially means drawing a conclusion without adequate support of reason, and often requires one to choose from possible alternatives or even, at times, to generate a possible choice where, apparently, there is none.

Guessing is so commonplace an occurrence in our everyday activities and our developmental process right from birth, continuing up to the terminal stages of life, that we are mostly unaware of its ubiquity and all-pervading importance. We keep on drawing conclusions based on guesses so continuously that most of the time it does not seem to involve *effort*. By contrast, *deducing* by following the path of regimented logic requires a great deal of effort, and engages our admiration. The great relevance of guessing goes unacknowledged in our everyday activity, and more so in science, that highly specialized form of inquiry and inference.

Your three-month old baby has become irritable and is throwing tantrums. The baby's father insists that she is hungry and is feeling neglected, and, on the face of it, he may not be wrong. But you assert that the baby is having sore throat and call in the physician. Both of you have your 'reasons' for your respective assertions, but none is conclusive in itself.

A football coach, on looking at the medical report of a player in his team suffering from an ankle injury for the last two months, selects him for the next important match. On being criticized for the selection he maintains that he has had a hunch that the player would make a magical contribution in the upcoming match.

A child, barely four years of age, was shown the drawings of two circles, one large and one small, and also of two squares and two triangles of different sizes. On then being presented with the drawing of an oval-like figure and a three-sided figure with curved sides, she identified the former as a circle, and the latter as a triangle.

All these are instances of inductive inference — a form of inferential activity whose instances span a stupendous spectrum from the most commonplace to the most creative and sublime.

“Philosophers since Hume have struggled with the logical problem of induction, but children solve an even more difficult task — the practical problem of induction. Children somehow manage to learn concepts, categories, and word meanings, and all on the basis of a set of examples that seems hopelessly inadequate. The practical problem of induction does not disappear with adolescence: adults face it every day whenever they make any attempt to predict an uncertain outcome. Inductive inference is a fundamental part of everyday life, and for cognitive scientists, a fundamental phenomenon of human learning and reasoning in need of computational explanation.” [117]

Inductive inference will be of central relevance in this book, which is why I will start by briefly explaining a number of basic ideas relating to induction. First, the business of *defining* induction. What, specifically, are the characteristic features that define induction? Here I cannot give you a universally accepted definition since people distinguish between different types of induction and adopt different points of view in characterizing induction. A commonly adopted approach is to counterpose it against *deduction*. In other words, one classifies all inferential activities into two major classes, namely, induction and deduction, and then makes further divisions within the class of inductive inferences. However, neat and clear-cut classifications are not always possible, and one

quickly finds that things get mixed up to no small degree. In particular, one finds that the way deduction and induction appear to be related to each other depends on whether one adopts the *problem view* or the *process view* [57].

The problem view and the process view

The problem view distinguishes between deductive and inductive arguments on the basis of *form*, without overt reference to content. Consider, for instance, the following:

when I put on my brown coat I look like a joker / I have put on the brown coat today / I must be looking like a joker,

which is a deductive argument since its form guarantees that if the premises are true, then the conclusion will necessarily be true. It doesn't really matter whether the premises are actually true (the first proposition is definitely not true, I assure you — I don't *have* a brown coat).

Consider, on the other hand, the following:

I feel like an errant child whenever my Aunt Agatha is around / I am feeling like an errant child now / Aunt Agatha must be around somewhere,

where the form of the argument does not guarantee the truth of the conclusion even when the premises happen to be true (my maths tutor of bygone days happens to have a similar sinister influence over me). In this particular example, the argument happens to be an instance of *abduction* — the formation of a hypothesis on the basis of evidence which, according to some, constitutes a type of inference that is, to some extent, distinct from induction proper, the latter being an argument that constitutes a *generalization*. Thus, consider,

all dogs have tails / dogs are mammals / all mammals have tails;

which, manifestly, is a generalization (indeed, one too much so) and qualifies as an

instance of induction. Both of the last two arguments have one thing in common — there could be *alternative* conclusions compatible with the truth of the premises (other great personalities such as my maths tutor that make me feel like an errant child, and, some mammals not having tails). *This* is what will, in this book, be taken to be the distinctive feature of induction. To be sure, an inductive inference does have other features too that turn out to be relevant in various other contexts. But an inference in which the conclusion is not a *necessary* one, and other similarly contingent conclusions are also possible, defines a class that is of overwhelming prevalence and relevance.

If the problem view identifies inductive inference in what can be termed an objective manner, distinguishing it clearly from deductive inference, the *process* view is a subjective one, and refers to the psychological processes that generate induction. And, in this process view, the distinction between deduction and induction is not so clear. But, it is the process view that is of greater relevance when we ask the question as to how people actually perform inferential acts of an inductive nature. However, there is no single accepted account of the psychological processes underlying deduction and induction. It is possible that both types of inference are based on the same kind of processing while, in contrast, a *two-process* account is also of considerable explanatory power.

“According to one-process accounts, the same kind of processing underlies both induction and deduction. Another way to describe this idea is that there is essentially one kind of reasoning, which may be applied to a variety of problems that could be considered either inductive or deductive in nature In contrast, according to two-process accounts, there are two distinct kinds of reasoning. It is possible that these two kinds of reasoning directly correspond to induction and deduction. Alternately, the two kinds of reasoning might correspond to some other distinction, such as intuitive reasoning versus deliberative reasoning that could be related to the distinction between induction and deduction.” [57].

We will, in this book, be specifically interested in the processes underlying inference making in general and inductive inference in particular.

Processes of an inductive nature are commonly found to occur in organisms ranging from animals to human beings, and are now performed, in various manners and up to various degrees of complexity, by computer programs as well, where all these taken together are at times referred to as *cognitive systems*. However, in speaking of cognitive systems, I will mostly have *human* inferential activity in mind while cognitive systems of other descriptions will be, at times, also referred to.

Imagine that a cognitive system receives certain inputs from the external world and becomes activated so as to attain a certain goal. Or, it may even be the case that the 'input' is generated mostly internally (*a child suddenly feeling that she has not seen her mother for some stretch of time*), thereby setting the desired goal (*mother has to be located*) or, to put it differently, setting a 'problem' to be solved (*a husband receiving a divorce notice and frantically thinking as to how to revive his marriage*). It may quite conceivably be the case that the inputs are not ones that have routinely been received in the past and that the goal is also not one routinely faced by the cognitive system, because in the case of routine inputs and routine goals, the system may make use of certain *rules* that have been learned in the past in solving the relevant problems, where these rules are 'objective' in the sense that other cognitive systems, similarly endowed, would also have made use of similar rules.

Inferential activities are set in action when inputs (generated internally or externally, or, in parts, both internally and externally) make it necessary that a problem (in a general sense) be solved or some goal be attained. If the inferential process is based on rules learned in the past that are not specific to the cognitive system under consideration and that, when applied appropriately, are guaranteed, more or less, to generate a correct solution to the problem at hand, then one has a case of deductive inference. On the other hand, if the inputs and the goal have novel elements in them (*mother absent for a long time and everything too quiet all around*) then one has to generate a solution all by oneself, and for this one looks for clues, again, in part from within (from past experience stored in memory) and in part from the external world (clues, additional information, overall context). One then generates, implicitly or explicitly, a *hypothesis* that points a way to the attainment of the goal (it crying out loudly, and crawling out through the

door toward the kitchen where mother was found on a previous occasion), a hypothesis that would differ to a greater or lesser extent from one cognitive system to another and one that may even lead to a wrong solution to the problem at hand (*mother not in the kitchen*).

Evidently, no idealized description is possible in respect of such an inferential process, which makes a clear distinction between deductive and inductive inference not an easy thing to achieve. We do have very definite instances of *deduction*, though, namely, *mathematical reasoning*, based, in the ultimate analysis, on the *rules of logic* (and, in addition, on the postulates and rules of *set theory*). A child is taught the rules of adding numbers, and she makes use of these rules whenever asked to work on an addition problem. A friend of hers also makes use of the *same* rules, and *both* arrive at the correct solution to the addition problem. But then, mathematical reasoning tasks can be stupendously complex and difficult as well, where it is not simply the matter of a sequential application of known rules and arriving at the correct answer. Quite frequently, one does not know *which* rules to apply in what sequence so as to arrive at the solution of a mathematical problem, and one has to *decide* on a course as to which route to follow in order that the correct solution be found, which needs inspired guessing an essential component of inductive inference.

“It is most unlikely that more than a tiny minority of mathematical theorems were ever in fact arrived at, “discovered”, merely by the exercise of deductive reasoning. Most of them entered the mind by processes of the kind vaguely called “intuitive”; deduction or logical derivation came later, to justify or falsify what was in the first place an “inspiration” or an intuitive belief.”, ([88], p 26).

“Many researchers in the field of philosophy, logic, and cognitive science have sustained that deductive reasoning also consists in the employment of logical rules in a heuristic manner, even maintaining the truth preserving character: the application of the rules is organized in a way that is able to recommend a particular course of actions instead of another one.”, ([82], p 48).

George Polya made remarkable contributions to the analysis and understanding of mathematical reasoning, and stressed upon the role of inductive inference in all kinds

of creative mathematical thought (recall, from chapter 3, how Polanyi acknowledged Polyas work on the role of heuristics in mathematical explorations; heuristics have a great role to play in inductive inference, as we will see later in this chapter, and in chapter 6). Polya underlined the non-deductive character of creative mathematical thought by making use of the phrase ‘plausible reasoning’ [98], contrasting it with ‘demonstrative reasoning’, the latter serving essentially the purpose of justifying the inferences arrived at by plausible reasoning. Plausible reasoning is another name of informed and inspired guessing. It constitutes a principal component of the psychological process of discovering mathematical truth. Commonly, a mathematical proposition is first guessed at and subsequently proved with considerable effort. At times, the proposition resists all attempts at proof and, at the same time, cannot be dismissed as being false, a celebrated example of such a proposition being *Goldbach’s conjecture*.

However, it is always the way that mathematicians and scientists make use of well defined *rules of logic* that captures our imagination as the quintessential feature of reasoning and inference. It is the set of these rules (along with the postulates and rules of set theory, forming the secure basis of much of mathematics) that has, in common perception as also in the perception of philosophers, set the *standard or norm* of objective reasoning and inference. The logical school of philosophy of science sets great value to the logical structure of scientific theories but, in real life, theories are seldom constructed in a logically tight structure (Robert Klee, in [73], chapter 2, gives an instance of how a theory might look when formulated in such a fashion). Instead, scientists take great pains to ensure that the theories are not *inconsistent* with known concepts, results, and evidence of proven worth. In practice, the ideal of objective reasoning following the rules of logic constitutes just one extreme end of an enormously wide spectrum of inferential processes, and only the instances belonging to this extreme endpoint truly deserve the name of deductive reasoning. In principle, anything within this wide spectrum, not coinciding with the endpoint of rigorous deductive reasoning falls within the ambit of inductive inference, which is why an overwhelmingly large fraction of all real life inferential activity has to be identified as being inductive in nature.

The enormous importance of deductive reasoning in everyday experience and in scientific explorations cannot, of course, be underestimated. What has to be scaled *up* is our level of acknowledgment of the relevance of inductive inference. Deductive inference is always carried out in the reasoning process of individuals and groups for justifying hypotheses and inferences arrived at in an inductive manner. However, inductive inference and its justification mostly proceed hand in hand, and the two can be distinguished only notionally. Speaking in relative terms, *pure* deduction is rare indeed.

That the problem view and the process view distinguishing between deductive and inductive inferential activity by cognitive systems do not correspond to each other is demonstrated by means of *psychological experiments* where, typically, a psychologist poses a problem to the subject that, on paper, appears to be one of deductive nature. In other words, the subject is required to perform an act of deductive reasoning according to the *problem view* — one that can be executed by appropriately following and making use of the rules of logic known to him (the subject). What is commonly observed in such experiments is something else — the subject works out a ‘solution’ that does not routinely agree with what the rules of logic would decree, i.e., in other words, the subject follows a *process* that does not necessarily correspond to the problem view.

In summary, inductive inference belongs to a wide spectrum of processes, of which only one extreme corresponds to truly deductive reasoning which is rule-based, the rules being independent of the particular cognitive system doing the reasoning, and guaranteeing to produce the correct solution to the problem at hand. The ideal prototype of deductive reasoning is provided by the rules of logic which, when applied sequentially along with a number of basic rules of mathematics (those of set theory, to be precise) do lead the cognitive system to the correct solution.

Deductive reasoning is essentially formal in that, given a set of premises, the rules operate on these without regard to their content or meaning. In other words, deduction is *syntactic* rather than *semantic*. The latter involves the meaning and relevance of concepts and propositions in some larger *context*.

Inductive inference: contextual, ambiguous, and non-deductive

Inductive inference is, in contrast to deductive reasoning, *context-dependent* where the *meaning* or the significance of a set of initial premises or inputs is evaluated within a given context and then a process is initiated that is more or less specific to the cognitive system doing the inferential act (for instance, the subject of a psychological experiment) and where the factors — of whatever nature — driving the process are not laid out with any degree of clarity to be described and interpreted by external cognitive systems (the psychologist). Inductive inference is fallible and uncertain — where the inductive process results in a conclusion which is chosen from among possible alternatives within the relevant context. The ‘choice’ is once again context dependent and may, in fact, turn out to be an erroneous one when viewed against the goal set out to be arrived at.

Put differently, inductive inference includes a wide spectrum that can be referred to as *non-deductive*, where pure deduction corresponds to an extreme end of the spectrum and is in the nature of an exception. Non-deductive inference is, mostly, of a mixed type where context-independent rules have some role to play but where, at the same time, complex context-dependent processes assume relevance (processes, moreover, specifically dependent on the particular cognitive system under consideration — our ‘system’ of interest being an individual person). The meaning of the term ‘context’ will emerge in the course of the remainder of this book.

We will, in this book, try to have a look into the type of ‘processes’ that are involved in the inductive inference of an individual, thereby connecting with the *personal* aspect of knowledge that Michael Polanyi spoke of. What is important to take note of here is that a non-deductive inference, not being rule-driven like a deductive inference, involves *gaps* that are not bridged by the application of rules akin to logical ones. The individual making the inductive inference performs, in a manner of speaking, logical *leaps* in crossing these gaps, and it is precisely the necessity of these leaps that results in what is *specific to the individual* person making inference, and that engenders the possibility

of alternative inferences in the process — some other individual would perform the leaps in some different manner arriving, possibly, at some other conclusion consistent with the inputs she started from (*two experienced physicians, separately examining a patient meticulously and thoroughly, and then arriving at different diagnoses based on the clinical symptoms*) or, in some case may even *fail* to effect a crucial leap or two.

Before we proceed, I want to repeat that when we speak of a ‘non-deductive’ inference in this book, we will refer, in a general way, to an inductive inference (the apparent circularity in nomenclature notwithstanding). For me, the crucial feature of inductive inference is the possibility of alternative conclusions consistent with the inputs initiating the process of inference and the consequent necessity of making one or more choices that result in some particular conclusion to the exclusion of other possible ones, these choices not being dictated unambiguously by strict logic. At times, the term induction is used in the more specific sense of an inference that produces a generalization such as the following:

the two items that I was served with in this restaurant were delicious / all the items available in this restaurant must be delicious.

There are, however, other types of non-deductive inference that differ from inductive generalization, while being intimately related to it.

In the remaining part of this chapter, I will first come back to the question of a workable definition of induction, along with a brief survey of the taxonomy of induction — what various different *types* of inductive processes there are. The latter will serve the purpose of definition by enumeration, and will make for a more focused characterization of induction, highlighting the overwhelming importance of induction in human cognition and behavior. I will, in this context, briefly explain the idea of *abduction*, the formation of explanatory hypotheses — an inferential process of great relevance in science. While abduction is considered by some as a special type of induction, with induction defined in a general sense, others are in favor of looking at abduction and induction as two distinct, though closely related, cognitive processes (refer to [82], [40], and to chapter 8

of this book for a detailed examination of the process of abduction). I will then present a brief introduction to various factors of relevance in explaining the process of induction, which will subsequently help us address the issue of human rationality. Finally, I will close this chapter by, emphasizing on the role of intrinsic factors in the inductive process, and their context-dependence.

Induction: definition, features, and taxonomy

Holland et al ([61], p 1) define induction as a process that “expands knowledge in the face of uncertainty”. Here ‘knowledge’ means things learnt in the past or, more generally — and a bit more vaguely — things believed to be true from past experience. And ‘uncertainty’ refers to the fact that the input data, i.e., the facts of current experience that set the act of inference in motion, are not sufficient to uniquely lead to a conclusion — the latter being the premise(s) constituting the ‘solution’ to the problem at hand.

Consider once again the imagined experiment with the child who was shown drawings of geometrical figures of various shapes and sizes and was told as to which ones were circles, which ones triangles, and which ones squares. These instructions get stored in her memory as ‘knowledge’ gained from past experience. But note that she was given only certain examples, and not precise definitions of the various geometrical shapes, since the terms of the definitions would not be meaningful to her. On then being presented with an oval-like figure and a triangle-like shape with curved boundaries and corners, she is found to identify these as a circle and a triangle respectively.

Neither of these identifications could, in the literal sense, be correct, and she could very well offer a confused response, which would be quite in keeping with the nature of the problem presented to her. Yet she did respond (with a charmingly majestic confidence) in each case, and the response did her credit too since she did identify some features that were common to the drawings previously shown to her and the ones she was asked to identify. A figure with three curved boundaries and curved corners could conceivably be identified as circle-like, and the oval-like figure could conceivably be identified as triangle-like. But the child did come out with identifications in the face of uncertainty

and, on being told that her identifications were on the right track, the entire process did expand her knowledge of 'circle-like' and 'triangle-like' figures. This 'knowledge', which is admittedly not precise, or 'correct' in the literal sense, would later provide the context in which she would learn and make sense of precise definitions of a circle and a triangle, but the ideas of 'circle-like' and 'triangle-like' would continue to stand for valid ones for her, along with many more instances that she would doubtless come across in course of time.

The noted psychologist, Philip Johnson-Laird has had the following revealing comment to make on induction ([68], p 169):

"In fact, much of our reasoning is inductive and outside the scope of logic. Our conclusions may be true, but even with true premises no guarantee can exist for their truth, because induction is fallible. I mentioned that textbooks often define induction as reasoning from the particular to the general. But ... the definition isn't quite right. We can make inductions that make particular conclusions. And so I defended this working definition: inductions go beyond the information given, and rule out more possibilities than their premises do. That is their hallmark."

In this paragraph, Johnson-laird refuses to define induction simply as an inferential process resulting in a generalization since inductive inference, at times, may involve the identification of a particular member as belonging to a certain class (identifying a face on having been given a number of characteristic features such as blue eyes, thick lips, and a sharp nose), or processes of other descriptions as well. On the contrary, what he sees as the hallmark of induction is that it leads to conclusions (or actions based on conclusions generated internally, as in decision making) that do not *necessarily* follow from the premises, conclusions *consistent* with those premises, but ones arrived at after *ruling out* some possible alternatives. This, in other words, involves a *choice*, where the choice is not determined by the premises *alone*, but is arrived at by means of a 'leap' across a void left open by the fact that the premises do not provide any key as they do in a deductive inference where the subject moves through a succession of intermediate stages by following a set of well defined rules (recall those uneasy days

when your maths tutor tried to drive home to you the derivation of those abominable theorems of Euclids). We will return later to the question of how the choice is made, what the process of making an inductive inference consists of. What is of relevance now is to have a feel of what is involved in induction — an idea, even if a vague one, of the ‘hallmark’ of induction, as Johnson-Laird puts it.

Another way to highlight the basic features of induction is to say that induction increases the *semantic* information in the premises, where the term ‘semantic’ refers to the meaning residing in the premises within the context in which the inferential activity takes place, and ‘information’ to the extent to which possible alternatives are ruled out in making that meaning more concrete, i.e., in seeking to *add* something to the meaning. For instance, when we are given clues for identifying a face (blue eyes, thick lips, sharp nose) and we do identify a particular face on the basis of those clues, the terms like ‘thick lips’ and ‘sharp nose’ receive an added texture of meaning (in virtue of the specific identification), and information is gained in that other similar faces are ruled out in making the choice (*thick lips all right, but eyes not so blue, and nose a bit bent*). Here is another instance of how information is gained by the ruling out of alternatives: if there are five students in a class then the knowledge that a drawing has been done by one particular student carries some information; but if there are ten students, then the same piece of knowledge carries more information because it eliminates nine possibilities while in the former case the name eliminates only four possibilities. Here the gain in semantic information (carried by the name of one particular student among all the students in the class — one who did the drawing) increases as more possibilities are ruled out.

Based on these ideas relating to the basic features of inductive inference, the following may be identified as a few of the more conspicuous types of induction that we routinely undertake in our everyday activities as also in higher cognitive activities such as reading and interpreting a highly contentious essay on the international political situation, getting to the bottom of a sublime piece of poetry, or working through the steps of a complex experiment in the chemistry lab.

1. *Generalization.* This is the most common type of inductive inference mentioned in texts. Suppose you have come across a number of objects of a certain type, all sharing a common characteristic. You then infer that *all* objects of that type share that same characteristic. For instance, on examining a number of items in a vegetable shop newly set up in your locality, you find those to be not of good quality. You then warn your friend that the shop is not to be patronized since it sells vegetables of inferior quality.

2. *Category formation.* A child is shown a particular car for the first time in her life and is told that the object is a car. On then looking at another car she immediately shouts in joy that she has seen another car. In this case, the child has picked out on certain features of the object that was identified as a car to her — features that appeared to be salient to her senses (for instance, the wheels), and ignored other features that did not appear to be as remarkable. On next looking at an object that possesses those salient features, she blurts out that it is a car, thereby arriving at the suggestion of a category — that of cars, even though there are, in fact, a thousand and one differences in particular features of the two cars that she has seen. A related activity is that of identifying subdivisions within a known category. For instance, on being told of certain characteristic features of two breeds of dogs, and then seeing a dog across the street, a child correctly identifies its breed. Here she focuses on the distinguishing features of the dog in question from the features of the other breed as told to her, and ignores the similarities, perhaps less salient to her (in the context of identifying the breed), with the latter.

3. *Analogical reasoning.* A nineteenth century scientist, on experimenting with the signals emitted by a radio transmitter, has a sudden feeling that the signals have certain similarities with light rays. Subsequently, optical signals and radio signals are identified as two types of electromagnetic waves, belonging to two different ranges of wavelength. Identifying analogies across

apparently unrelated categories is an inferential act of remarkable relevance in everyday activities as also in scientific discovery, engendering the possibility of a supremely creative act.

4. *Making judgments and decisions.* As an examiner in an essay competition, you are at a loss while judging the relative merits of two particular students. While one of the two is endowed with a decidedly better literary flair, the depth of understanding of the other one is really remarkable. Which of the two would you judge to be the more deserving candidate to receive the first prize in the competition? You are at a loss because the two qualities — one of literary flair and other of depth of understanding — are disparate ones. The judgment that you finally make is crucially dependent on values specific to you and not determined solely by features or qualities of the two essays in question. Some other examiner could very well have made a different judgment, in which case the prize would go to the other candidate.

Decision making is another type of inductive activity closely related to making a judgement.

“In order to decide, judge; in order to judge, reason; in order to reason, decide [how reason]” [69].

While teaching a class of teenagers, you have a feeling that one particular student has a behavioral issue. You observe her keenly for as long as one month, comparing her traits with those of her class-mates, and find yourself in two minds as to what course to adopt regarding the child. Finally, you call up her parents and suggest that they consult expert counsel, even as you have the nagging doubt that the child may, after all, be having no problem at all, being just different from the rest. Your decision rests on your judgment that the child is not simply different from the rest, but is having some kind of a problem in her social interactions.

5. *Abduction*. Abduction is the inferential process of making hypotheses that have an explanatory power and is of great relevance in, among other things, *scientific thinking*. Science seeks to explain, to build new concepts and theories that expand existing theoretical frameworks, when the latter fail to explain facts of observation or to make a coherent whole of the facts and individual concepts making up the framework. At times, auxiliary concepts and theories are postulated as adjuncts to an existing theoretical framework so as to consistently accommodate known facts and available concepts within it, without substantial alteration of the latter. Once in a while, an entirely novel hypothesis is put forward that subsequently transforms an entire conceptual framework. In all this, the scientist makes a *conjecture*, in the nature of a logical leap that subsequently passes through a process of appraisal, partly in the unconscious mind of the scientist herself, partly in a deliberate experimental and logical process of appraisal undertaken by her and then, in the remaining part, through the experimental and theoretical scrutiny of her peers in the scientific community.

Charles Sanders Peirce, the American philosopher, is credited with a substantial analysis of the abduction process of inference. Peirce's ideas on abduction were formulated in stages, where, in the early stages Peirce was mainly concerned with the *syllogistic* theory of abduction — a theory of looking for one or more ‘missing premises’ in a logical argument where the conclusions are known in part or in whole in the form of puzzling statements or observations that the current body of knowledge (again, in the form of a number of premises) fails to lead to. Abduction, in this view, is ‘reasoning in reverse’, which implies that the premises arrived at by abduction are of uncertain validity since the relation of logical implication is not a reversible or reflective one (‘A implies B’ does not imply ‘B implies A’). Related to this view is the subsequent idea of Peirce where he advanced the inferential theory of abduction, in which “abduction represents the hypothesis generation part of explanatory reasoning” ([40], p5).

Peirce distinguished between induction and abduction while, at the same time, leaving the door open for subsuming both under a broader description of non-deductive reasoning. Lorenzo Magnani points out that a pervasive view in philosophy, of the process of hypothesis generation, looks at it as one that is “paradoxical, either illusory or obscure, implicit, and not analyzable” ([82], p 1; see also p 21, 27 where the term ‘retroduction’ has been introduced, in the same sense as ‘abduction’). It may be noted, however, that certain instances of induction and abduction have been generated in artificial cognitive systems (see [40], [82]).

The process of abduction will be taken up at greater length in chapter 8.

6. *Providing support and confirmation*, in a limited sense, to generalizations and hypotheses from observed facts. This brand of induction has evidently a considerable overlap with induction as the inferential process of generalization, but is possessed of a distinct aspect as well, where facts of observation are made use of in inferring the validity of hypotheses and theories. Indeed, induction has been looked at in the philosophy of science as providing the *logic* of confirmation of scientific hypotheses and theories, which, in a sense, is antithetical to the view that induction is an inferential process of an essentially non-logical and fallible nature. I will have more to say on this by way of explanation in paragraphs below, where it will be seen that the program, in the philosophy of science, of looking at induction as a logic of confirmation did not quite meet with success, and induction was instead accepted as an inferential process that can provide support, in a limited sense, to hypotheses and theories. This role, however, is distinct from abduction, which is the process of *generation* of hypotheses.

Now that we have some idea as to the defining features of induction as non-deductive inference, and to a number of types of induction, where the latter includes the process of abduction — an inferential process of great relevance in science — I will raise a number

of issues that will set the stage for subsequent chapters of this book. However, before I do that I must mention that induction, as outlined above, covers almost the entire range of mental activities in human *cognition*, where cognition, broadly speaking, refers to the aspect of information processing in the human mind, such processing being, in general, in the nature of a sequential activity that may or may not be a well defined or explicit one. Cognition, in other words, may involve a *sporadic* succession of steps rather than a clearly defined sequence based on definite rules, where the steps may be interlaced with tangled heaps of branches, possibly having unconscious moorings. While ‘cognition’ is a term having a broad coverage, ‘induction’ is one with a slightly more specific connotation — that of being an *inferential* activity in relation to a goal or purpose, set in motion by inputs from without or within. Induction, indeed, draws upon almost all mental processes associated with cognition and, conversely, is involved in some form or other in almost all such processes.

Before I proceed, I briefly deal, in the next section, with the question of whether and to what extent induction can provide us with a logic of confirmation, because this issue has been the cause of a vast body of work and polemics in the philosophy of science. Summarily stated, the very fallibility inherent in the inductive process goes against the possibility of induction providing a secure foundation for the confirmation of scientific hypotheses.

Can induction constitute a logic of confirmation?

At the outset, I have to tell you that the terms ‘induction’, ‘inductive logic’, and ‘inductivism’ are used in two different contexts that, paradoxically, are of contrary significance, at least on the face of it. As I have indicated above, a number of philosophers of science have used these terms to look for and to describe a *logic of confirmation* of explanatory hypotheses, hoping that induction may thus be given a place analogous to deductive logic, especially in scientific explorations. In scientific inquiry, people continually make *hypotheses* to explain numerous facts of observation where the psychological process of arriving at these hypotheses is obscure, to say the least. However, taking for granted one

or more such hypotheses in the context of a set of observed facts, can one find a *logic* whereby one particular hypothesis, which may be made up of a number of concepts, and which may assume the form of a theory, is confirmed by these facts in preference to other possible, alternative hypotheses? If so, then this logic of confirmation of hypotheses from an analysis, scrutiny, and collation of facts of observation can be given the name of ‘inductive logic’, which will then constitute the standard or *norm* for testing and comparing scientific hypotheses and theories in an objective and impersonal manner.

It is important to note that such a purported norm has to be objective, or *rational*, in the dual sense of being impersonal, i.e., free of the vagaries of the psychology of individuals, and of being capable of pointing out the *true* significance of observed facts relating to *actual* properties or qualities of natural objects.

As opposed to such a logic of confirmation, the term ‘inductive inference’ will be used in this book to stand for the inferential process of arriving at generalizations, hypotheses, and *fallible* conclusions that do not, in general, conform to desired norms of objectivity and rationality. On the one hand, these are in the nature of *interpretations* of qualities, properties, and correlations existing in natural objects while, on the other, the process of generation of these interpretations may have deeply entrenched roots in the minds of individuals.

Ladyman ([78], sections 1.3, 1.4) refers to two senses in which the term induction can be used when he speaks of the ‘new tool’ (*Novum Organum*) of induction proposed by Francis Bacon on the eve of the scientific revolution in the western world:

“Induction in the broadest sense is just any form of reasoning that is not deductive, but in the narrower sense that Bacon uses it, it is the form of reasoning where we generalize from a whole collection of particular instances to a general conclusion”.

Here the ‘narrower’ of the two senses points to the approach of looking at induction as providing a possible logic of confirmation of theories from facts of observation. However,

as I have mentioned above, the support that induction can provide to a hypothesis or a theory can at best be a limited one because of the inherent fallibility of inductive generalization which is essentially an interpretation of observed facts rather than a logical confirmation of a hypothesis on the basis of those facts. On the one hand, philosophers would like to see induction as providing a rational norm for effecting an expansion of knowledge in an inferential process where one reaches beyond the premises that one starts from while, on the other, a psychological examination of the process of inductive inference fails to identify such normative and rational hallmarks in induction.

The proposed use of the term ‘induction’ in the sense a logic of confirmation did not quite find a secure ground for itself, even when its intended meaning is broadened to include probabilistic confirmation, where a set of observed facts is made use of to infer whether the validity of some particular hypothesis or theory is more probable than that of others (refer to [78], chapter 2, [46], chapter 2, chapter 14). David Hume, for one, successfully propounded the view that there is ground for profound skepticism to induction as the logic of confirmation, and that inductive inference really has no ‘logic’ of the same standing as deductive logic. In other words, induction, properly speaking, is a non-deductive mode of inference that does not conform to standards of *objectivity*.

There are, however, subtleties here. Once one accepts that there is no justification in using induction as a logic of confirmation ([78], chapter 2) endowed with the credentials of objectivity, one is led to looking at induction as non-deductive inference in a *broad* sense, as pointed out by Ladyman. And, as I have indicated above, within this broad connotation, there exists room for the view that induction, in a narrower sense, does have a role to play in the *acceptance*, if not in the confirmation, of hypotheses and theories. Suppose we have a theory A that is consistent with a set of observed facts and with proposition(s) deriving from some other hypotheses that have been justified on more or less secure grounds (there is, of course, no *absolute* justification for *anything*). One then does use this as a ground for acceptance of the theory A, at least provisionally. This is the sense in which one agrees with what Ladyman means when he speaks of the term ‘induction’ as having two different connotations one broader and the other narrower.

Indeed, there exist different points of view regarding the possible role of induction as providing, in a sense, a logic of confirmation of hypotheses that are now centered around the search for a *probabilistic* justification of induction. The latter is a refined and attractive version of the earlier, relatively naive, view that a painstaking program of collection and examination of facts of observation in a given domain of discourse enables one to choose a hypothesis that provides the best explanation of the facts and, if the process of referral to facts is carried forward appropriately, can even provide a ground for confirmation of that hypothesis. There can be found an extensive literature on the *Bayesian theory* [46], [116] of justifying an inductive inference, one that is supposed to promise *normative* appraisal of such inferences, analogous, in some sense, to the normative role of deductive logic in respect of inferential reasoning acts belonging to a certain class well known in exact sciences (indeed, all mathematical derivations may be said to belong to this class). The Bayesian theory has found applications in *artificial intelligence* where particular tasks are carried out having features common with human inductive inference. This raises the question as to whether there is a probabilistic basis of human reasoning and inferential processes in general, and how far this can provide the ground of a rational appraisal of an inferential act. This we will briefly come back to in a later section in this book (refer to chapter 7, section entitled *The rationality issue: a brief overview*).

Induction: questions and issues

With this much of an introduction to the idea of inductive inference, I will now raise, as promised, a number of issues that will enable us to gain a deeper understanding of induction — what it stands for and how it is to be viewed in a broader perspective of reasoning and inference making in the context of human cognition in general. Some of these can be looked upon as pointers to a theory of induction where aspects of human cognition are examined and analyzed, telling us how induction can possibly be realized in the human mind. A number of these issues will be addressed in greater details in subsequent chapters of this book.

The term ‘theory’ in the field of cognitive science does not carry quite the same connotation as that applying to theories of physics or chemistry since theories in cognitive science are relatively more speculative and vague, serving only to provide us with plausible explanations or viewpoints in respect of this or that area of cognitive activity. Controlled experiments of substantial significance are rare, and definitive conclusions from experimental observations are equally rare — at times providing almost comparable support to alternative and competing theories. Indeed, these ‘theories’ are themselves in the nature of inductive inferences with quite wide gaps remaining in their structural and logical organization and in their relation to facts of observation. Still, these theories of human reasoning and cognition are useful and valuable pointers — ones at the cross-road of science and philosophy — towards an understanding of human cognition in general, and inductive inference in particular. From the philosophical point of view, theories of induction, sought to be arrived at from the experimental and theoretical base of psychology, provides an instance of the *naturalistic* approach in philosophy. This approach is likely to give us new insights relating to traditional problems on epistemology, to which analytic philosophy does not provide satisfactory answers since it adopts a contemplative approach rather than looking at how exactly we come to possess our beliefs and our knowledge about the world.

Here is a warning. In the following sections in this chapter, I will place before you a number of observations and statements some of which, *especially those in the section highlighting on beliefs and emotions do not enjoy direct support in the cognitive science literature* (the same goes for parts of chapters 6 and 7). But, at the same time that I warn you on this, I also want to assure you that these are not inconsistent with the literature either. I look at these as *plausible* ones — I wouldn’t know if you will agree with me. I put these here nevertheless as indicative of a possible framework for having a good understanding of the process of inductive inference which is, as of now, a baffling one indeed. Here I go.

Unconscious cognition

The first issue I want to raise is that of *tacit cognition* and the role of the unconscious in human cognition. The realization that *unconscious* processes play an all-pervasive

role in human cognition is a relatively recent one. Of course, the unconscious has been looked upon for centuries as a playground of mystical and dark forces rampaging in the human mind, but has mostly been confined to folklore, poetry, art, contemplative philosophy, and literary imagination, without a disciplined investigation into its role in cognition. It took center stage in the science of psychology in Freudian investigations into the human mind, but then got eclipsed as Freudian psychology itself fell from grace in the academics and the viewpoint of behaviorism attained a position of dominance, when even the word ‘unconscious’ was anathema in academic psychology. The unconscious once again found a place of its own by the mid-eighties of the last century in the shape of a cognitive unconscious which differs somewhat from the Freudian unconscious, though one cannot rule out the possibility of an appreciable overlap and a correlation between the two (see, for instance, [126], [16]).

Beginning from the late nineteen sixties, several streams of investigation became visible in cognitive psychology, with their orientation towards the cognitive unconscious. These investigations, while maintaining their distance from the psychoanalytic unconscious, quickly broadened in scope and richness.

Mankind has wondered about intuition — intelligence residing in the unconscious — for ages. The great physicist and philosopher Helmholtz highlighted the role of unconscious inferential processes in visual perception ([68], p52; see [103], p15, for references to early speculation on the cognitive unconscious) and of intuition in scientific creativity, citing the case of Michael Faraday as a prime example of the workings of the intuition ([11], p 36). The notion of an unconscious reasoning mechanism distinct from the deliberative reasoning capacity of the human mind, such as the capacity for mathematical reasoning, was put forward by the philosopher-scientist Pascal ([68], p 53) — a notion that anticipates the more recent *dual-process* theories of reasoning that we will have a look at later in chapter 7 of this book. And a good number of scientists, mathematicians, and poets, while introspecting about the source of their inspiration in creative contributions, have wondered about their own unconscious capacities.

As for the relatively recent resurgence of interest in unconscious cognitive mechanisms,

I am inclined to trace it back to the seminal contribution of Michael Polanyi who expounded on tacit knowledge as a separate storehouse of cognitive skills and abilities in the human mind (“a wholly explicit knowledge is unthinkable”; quoted in [23], p1). Polanyis ideas acted as a major inspiration for Arthur Reber who was among the first to systematically take up the study of these unconscious abilities in the course of theoretical and experimental investigations in cognitive psychology, and who was followed in quick succession by other pioneers, including Axel Cleeremans and John Kihlstrom who brought insight and diversity into the field. These investigations soon grew into a major activity, and terms like ‘implicit learning’, ‘tacit knowledge’, and the ‘cognitive unconscious’ soon gained wide currency, bringing to the fore the relevance and importance of unconscious cognition, which essentially relates to processes of acquisition, storage, retrieval, and manipulation of concepts, many of which are possibly of an atomic or elementary nature, without conscious intent, effort, or awareness: “the process by which knowledge about the rule-governed complexities of the stimulus environment are acquired independently of conscious attempts to do so” (Reber, quoted in [21], preface). Today, the investigations in unconscious cognitive activity cover a wider and much more diverse area compared to early days ([74], [56]) though, by the very nature of things, these involve a measure of speculative and indirect inquiry distinguishing it from other areas of scientific activity.

Though the range and variety of unconscious cognitive processes is thought to be enormously broad, we will be more specifically concerned with processes relating to inductive inference and to abduction (which is, broadly speaking, a type of induction, or at least a close cousin of it) where, in the context of the latter, we will be having a brief look at the process of *creative thinking*.

With more and more attention focusing on the cognitive unconscious, the question of rationality in human cognition and inferential activity has assumed great relevance. The ‘logic’ of the unconscious, whatever it is, is not the logic that we have a tendency to ascribe to the human mind. Are there identifiable principles governing the workings of the unconscious that we can hope to discover in days to come, when we will gain understanding of what now appears to be a pervasive lack of rationality in human rea-

soning and inference making? This question is, to a considerable degree, co-extensive with that of how inductive inference is carried out in the human mind, or, how exactly the guessing game of induction is played out.

Unconscious cognition will constitute the topic of chapter 5.

Induction: the role of heuristics

As we have seen, induction is essentially a guessing game, though one that is likely to have an underlying ‘logic’ not known to us. We need to guess in most situations in life where strict reasoning is either too difficult, too time consuming, or not in the realm of possibility due to paucity of information. Indeed, even when sound deductive reasoning is possible, we do not commonly make use of deductive rules in a consistent manner because, simply stated, inductive inference is our ‘second nature’. We guess because our inference-making is intrinsically opportunistic in character and has no in-built loyalty to quality or rigor of the reasoning process — what is more important than sound logic is *effectiveness*. This is a trait we have, perhaps, inherited in the course of our evolutionary history, and one that, moreover, manifests itself to various degrees in the developmental history of individuals.

One factor of great relevance in this guessing game is the use of *heuristics*. Heuristics are produced as partial and ready-made clues to the solution of a ‘problem’ that a cognitive system may happen to face (like, for instance, a baby trying to seek out her mother — *mother likely to be found in the kitchen*, or a chess player trying to find a good move at a critical juncture in a game — *grab the queen*), clues that are themselves in the nature of guesses, not usually arrived at by solid reasoning. For a person engaged in making an inductive inference, the heuristics help her along to reach a conclusion by making logical leaps, i.e., by making guesses. It is by making a judicious use of these partial solutions, or partial guesses, that the person climbs up, so to speak, an inferential staircase where, at each step, she makes use of heuristics activated at some lower level, so as to land on a higher level of her upward journey, finally arriving at the top. We are all familiar with ‘hunches’ or ‘gut feelings’ that help us to navigate in

uncertain and demanding situations, often propelling us towards a satisfactory solution to the problem at hand but, at times, leading us astray too.

Heuristics are used not only in inductive inferences but in deductive ones too where these appear, in a manner of speaking, as small ‘packets’ of truth residing in the cognitive mind from innumerable cues picked up from past experience, that find their use in making ‘fast and frugal’ inferences without, however, the attendant risk of making the inference fallible. In this context, see [104].

There exists an extensive literature of relatively recent origin on the issue of whether heuristics constitute a useful and essential means in our inferential activity, or are *defects* in our reasoning process, producing *biases* that make the reasoning deviate from norms of rationality. Do the heuristics constitute a ‘fast and frugal’ method in inference making, or are these, predominantly, possessed of a nuisance value in the context of the reasoning process?

It seems that at least some part of the debate over this issue is an exchange at cross purposes. Heuristics are indeed a great help in the pursuit of inductive inference while, *at the same time*, they constitute the source of fallibility in the inferential process which is seldom carried out by following explicitly formulated reasoning steps based on clearly formulated specific rules. The difference in the two points of view stems from a basic issue, namely, whether it makes sense to speak of *normative standards* of rationality in human reasoning, especially in inductive inference. The point of view that heuristics constitute an essential, if fallible, ingredient in inductive inference, rests on the recognition that, when one looks closely at the inferential process, one finds that, on the one hand, the course of inference includes segments or stretches that appear to be rule-driven and rational while, on the other, these stretches are interspersed with gaps that *cannot* be negotiated with such well-defined rules. In other words, normative standards *cannot* be set for the inferential process as a whole which contains *gaps* where, precisely, the heuristics come in.

We, in everyday interactions with our social and natural environment, incessantly keep

on using heuristics of various descriptions — *don't put all your eggs in the same basket* (quite reasonable, but not a guarantee of success), *bribery always works* (highly fallible, but still useful at times), *old is gold* (doubtful), *once you have decided on a course of action, never deviate from it* (headstrong) — all these are instances of heuristics that may help us take good and effective decisions in matters personal and social, though each is a half-baked belief, and may not work in a crunch situation. However, we often do make use of unproven beliefs like these in making judgments and decisions, and what is more, we are, at times, aware of which of our beliefs we are invoking and what judgments or decisions we are making. These are heuristics operating in the conscious domain. More importantly, heuristics may operate in the unconscious domain too. The unconscious mind is an ever-active system where elementary beliefs and concepts are born and continue to be operative, as participants of processes of an elementary nature, associating with other similar beliefs and rudimentary concepts, forming heuristics of a relatively more complex structure, and thus setting into motion a *hierarchical* process where heuristics act in conjunction with one another at various levels. It may so happen that this entire hierarchical process remains entirely confined within the unconscious mind, getting expressed in subtle ways in our *behavior*. Or again, it may surface into awareness where it may end up being part of conscious reasoning. In his widely read book 'Gut Feelings' [44], Gerd Gigerenzer speaks of heuristics appearing as hunches and gut feelings that dominate much of our reasoning and decision making, and calls these 'the intelligence of the unconscious'.

The *gaze* heuristic has been discussed to some considerable extent in the literature (see [44]; in this context see also [93] for a criticism of the idea of the gaze heuristic being an intuitive mechanism of a simple description). A fielder in a cricket match runs a big distance to catch a powerfully hit ball that makes a trajectory in the air. One can, in principle, work out the trajectory of the ball with mathematical precision provided that one has access to a great deal of data of a very intricate nature. But the fielder is no mathematician and has no computer at hand to get the calculations done — he, as a matter of fact, doesn't have the faintest idea as to what those calculations might be. But still, he makes the chase and eventually succeeds in taking the difficult catch.

This he does by the simple expedient of fixing his gaze on the ball and keeping his line of vision fixed with respect to his body. Does he do it with conscious intent, knowing whether and how such a tactic should work? Indeed, even if he is made aware of the trick by some colleague of his, much of his actual mental and motor activity occurs in the unconscious domain, involving a lot of processing, a massive use of cues, an equally massive aggregate of implicitly learnt lessons that have no more claim to truth than, perhaps, some success in the past, and a hierarchy of unconscious decisions, one feeding the next. Overall, the gaze heuristic seems to be partly unconscious and partly conscious.

A heuristic is a kind of belief, of an elementary nature or, maybe, of a more complex type (*when you are in the end game, play the king*). It is some kind of a rule, a rule of thumb, that tells us how to proceed to an eventual inference or decision in a given context, where the context itself is, generally speaking, a greatly complex one if we try to get down to describe it in the minutest of details, so much so that an actual description of it may be a task that is computationally *intractable*. But the heuristic guides us nevertheless, drawing from the context and pointing the way to a judgment, a decision, or an action. And, inductive inferences typically involve a hierarchy of judgments and decisions ('actions', in the general sense of the term). The success, as also the fallibility, of induction rests on the fact that it typically proceeds through a maze of beliefs.

The role of heuristics in inference making — especially in inductive inference — will be repeatedly referred to in this book, especially in chapter 6, where it will be addressed in greater details (see [this section](#), and [this too](#)).

Induction: beliefs emotions and affects

Beliefs form a great and complex web in the mind of a person. Many of the beliefs hang together with other beliefs in the web, one set of beliefs supporting and often reinforcing another, while some others are less coherent and may even be outright inconsistent when judged against many other beliefs in the web. In other words, one harbors within one's mind a belief *system* with a rich, interconnected, and often strange or curious

structure where there results a complex and dynamic pattern of tensions, pulls and pushes among the various components of that structure.

What is more, beliefs are sheltered in the unconscious mind where they reside and interact with other small and large beliefs of various descriptions, defying basic norms of rationality to a great extent.

“The vast majority of beliefs, however, are not likely to be conscious or reportable, but instead simply taken as granted without reflection or awareness. Such beliefs may be inferred from a subject’s behavior, but otherwise remain unconscious and enacted [in] largely involuntary [manner].”, [24].

A mother believes one of her two children to be more intelligent than the other, based, perhaps, on a single long-forgotten incident in the past and refuses stubbornly to heed to more recent evidence to the contrary while, at the same time, never letting go of even a tiny piece of ‘evidence’ in favor of her own belief. Perhaps, too, her belief finds support from some other beliefs hiding in her subconscious mind, relating to bodily features (*blue eyes, sharp nose*) that she takes to be indicative of intelligence. Not that she is *happy* in her belief — on the contrary, she may be relatively more loving to the child she takes to be weaker in the mind. Many of one’s beliefs one is aware of are in the conscious mind only to some extent, being entrenched within the unconscious like the submerged portion of an iceberg.

There has been a lot of philosophical work on how belief compares with knowledge and, to what extent beliefs can be said to be true or false. Knowledge is supposed to be justified true belief, but questions of justification and truth are notoriously problematic ones. We will adopt here the ‘commonsense’ or the ‘folk psychological’ view of belief, truth, and knowledge, in which truth is seen not as an absolute but a context-dependent concept, and knowledge is seen as belief that has been accepted as true in a given context. Beliefs are constantly in a state of dynamic tension where these are confronted with new facts of observation, and with requirements of consistency and coherence with regard to one another. Some of the beliefs entertained by an individual or a community are shielded,

in a relative sense, from ongoing processes of confirmation and justification, and are held as sacred, once again in a relative sense. Others do not enjoy a similar degree of shielding and are subjected to processes of justification, correction and revision.

Depending on the degree of effort expended towards the justification of a belief in terms of consistency with other beliefs and of facts of observation, these can be imagined to be arranged in a scale where, at one end of the scale are the beliefs that are protected and guarded most zealously while at the other are those beliefs that are checked and counter-checked conscientiously, and are eventually accepted as 'knowledge'. The ones that are certified as knowledge are then spared further efforts at justification till there takes place a change of context where new yardsticks of justification and confirmation make their appearance so that a big chunk in the erstwhile web of beliefs undergoes a surgical operation and a new set of beliefs gain the certification of knowledge.

Looking at the set of beliefs that are not confirmed or justified as knowledge in a given context, some are actively defended against justification and count as *dogma* (beliefs that have gained the status of knowledge also constitute dogma, but in a different sense) while others are not so defended, and are in the nature of 'informed guesses'.

In other words, we have, roughly speaking, *three* sets of beliefs — the ones that are frozen into dogma, others that are in the nature of informed guesses and are subject to processes of confirmation and justification and, finally, those that have gained the status of knowledge. However, this is only a very partial description of the web of belief, to which one has to add the other aspect, namely, that of the *depth* of a belief. This relates to whether the belief, with all its associated psychological components, resides principally in the conscious mind or whether it has unconscious moorings, due to which it acquires a distinct significance.

Complementary to this *structural* view of the web of beliefs, is the *functional* view that informs us of the roles that the various beliefs play in the great laboratory of the mind. Stated briefly, our set of beliefs provides us with an exhaustive *map* to help us navigate in a highly complex, confusing, difficult, and uncertain world (or, in a world which, in

some sense, is *perceived* to be such). But the map is a very tentative one, made up in a long series of guesses, trials and errors during which it is put through an innumerable succession of small and big tests. It is corrected and patched up from time to time, and has been found to be reliable in respect of only a few landmarks (and that too, barring earthquakes!). As to the rest of its vast collection of indications, warnings and flagpoles, they are *not reliable*, but these are *all that we have*. We navigate with their help, we fail, we make a patchwork of improvement here and an improvement there which we are ourselves not sure of and, in the midst of all this, we succeed in making a number of good landings, some of those breathtakingly good. This process of navigation involves the making of deductive and inductive inferences — on the one hand, making good use of the map and, on the other, making continual revisions in the map itself.

The functional aspect of beliefs includes a *psychological* one wherein these make each of us face the world (and ourselves too!) as a unique individual, equipped with a unique collection of psychological traits — an infinite number of finely tuned defenses, stances, approaches, conflicts, weaknesses, and strengths. And it is here that beliefs and their complex dynamics are linked with *emotions, affects, feelings, and moods* — a number of factors that are, in small or large part, rooted in our *physiology*. This makes beliefs essentially connected with and woven into a vastly more complex psychosomatic system that is, in a large measure, beyond our conscious control, at least in the short run. The logic of inductive inference, if there is one, is to be sought by looking into this exquisitely complex arena where the drama of cognition unfolds.

While speaking of the role of beliefs in providing meanings and explanations to ideas and events, Connors and Halligan, in their valuable paper on a cognitive account of belief [24], make the following observation:

“Given that any search for meaning will largely depend on pre-existing beliefs and knowledge, the outcome is likely to be highly personal and idiosyncratic. Overarching narratives that are implicit in subjects’ pre-existing beliefs may be particularly influential in determining the outcome of the search. In addition, subjects may adopt particular attributional styles — habitual tendencies to explain events in certain ways

..... — whilst also relying on heuristics to save on cognitive effort Subjects' emotion and mood may also be relevant influences at this stage. Explanations may be selected because they are congruent with a prevailing emotion or dominant mood. Anxiety, for example, may foster explanations involving threat or danger, whereas happiness might prompt more benign explanations. Alternatively, explanations may be selected based on their affective consequences Explanations, for example, that offer certainty and comfort or maintain self-esteem and internal consistency are more likely to be selected over other explanations that do not provide these benefits, providing they are sufficiently plausible and can be rationalized. Motivation and emotion may constitute a particularly powerful determinant of evaluative beliefs ”.

Within the spectrum of loosely and solidly grounded beliefs are to be found — the *heuristics* (refer to the above paragraph from [24]). These are beliefs (guess-works, really) that are especially active in cognition or — to be more specific — in inductive inference and, like all other beliefs, are themselves the product of an ongoing process of cognitive inference — at times goal-directed and at times apparently without identifiable goals. Once formed, these are used as building blocks for further inferential activity, a good part of which is inductive and goal-directed in nature. It is in this sense that inductive inference can be said to be a *hierarchical* process — one stage of inference feeding into the next, where smaller guesses are made use of in making up broader guesses, and where there takes place a concomitant process of verification, confirmation, and justification all along, the latter being an equally fine-tuned and fantastic one, giving the process of inductive inference a unique flavor and efficacy.

And, much of this continuing process of guessing and verifying goes on within the unconscious mind, with the involvement of emotions, affects and feelings. While a huge store of heuristics of innumerable descriptions supplies the building blocks for inductive inference, these heuristics, considered all by themselves, are still not sufficient for explaining the basic act of guessing because of the fundamental fact that guessing involves a *decision* by means of which a choice is made between available alternatives. This is where the *context* comes in, in the form of background knowledge (prior expe-

rience, beliefs, heuristics, unconsciously stored cues), which fundamentally constrains the inductive process, providing a relatively small number of alternatives to choose from, and then arises the really crucial step — choosing from among these alternatives. Background knowledge is not of direct help here, because what one needs to compare in order to choose are *disparate* entities, with qualities that *cannot be reduced to a common denominator*.

You have to choose between two students, one of whom has a great literary flair, and the other a remarkable depth of comprehension, there being no other academic criterion that can help you fix your choice. You are at a loss what to decide. There is no logical criterion available on which to base your choice, because literary flair and depth of comprehension are disparate qualities. It is here that *extra-logical* factors come in and make themselves felt — factors that include affects, emotions, and feelings. One of the students belongs to a socially neglected and poor background, while the other to an affluent one. Perhaps you have some sympathy for the socially neglected and deprived people, believing them to be more deserving of your attention. Perhaps you yourself have risen from such a background, and your belief has an emotional counterpart, generating a pleasurable affect in you, without any conscious awareness, as you contemplate choosing the less fortunate of the two students, and — you make the choice, based on this *remote belief* of yours (relating to your natural sympathy for the less privileged) that has an emotional content. It is much like making use of a pole in order to take a vault — your emotions and affects act as your pole or as your *psychological springboard* in making a mental leap. This, however, raises a question: the use of psychological props such as beliefs, not directly relevant to the issue on which a decision is being sought, and ones such as emotions and affects, appears like having recourse to a coin toss in making a decision. Making decisions is of vital importance in reasoning and inference making. Can it then be said that reasoning and decision making are guided by events of an essentially random nature. *The answer has to be, No!* But arriving at the answer involves going to great depths. At such depths, however, *nothing is clear anymore*. Still, we will take a plunge into these uncharted depths in chapter 6 in a spirit of fun and adventure.

Inductive inference: summary

I will now *summarize* these few opening observations on inductive inference that I have brought up in this chapter, adding a few concluding comments for the sake of clarity. Inductive inference is an all-pervasive cognitive act where one starts from a set of inputs, and proceeds to solve a ‘problem’, where the urge to solve the problem arises from some sort of ‘trigger’ among the set of inputs (*baby feeling a pang of hunger, too quiet all round, something unusual, where is mother?*). The quest to arrive at a solution to the problem involves a cognitive journey where there takes place a sequential processing of ‘information’ in which, at each stage of processing, information is drawn from environmental inputs and from a vastly rich internal store generated in experience and prior acts of processing (in addition to a repertoire of innate capacities produced by evolutionary means — more on this later). The information, moreover, is not amenable to description in terms of a fixed ‘alphabet’ (i.e., a specific set of symbols) and possesses a syntactic as also a semantic content, the latter depending on the context in which the processing of the information occurs. In this, inferential processes taking place in the human mind differs from computational processing of a comparatively simple kind.

At some of the stages of information processing, rules learnt in prior experience are made use of, where these rules resulted in successful solutions to relevant problems in the past though, in spite of history of prior success, few of these rules are foolproof. In the remaining stages of the inferential process there arises the need to adopt a choice between possible alternatives where environmental cues and stored information are not of direct help, as a result of which there arises the necessity of a logical leap. And here, unknown to you, apparently remote beliefs (so remote as not be counted as relevant information in the context of your problem) linked with emotion and affect help you as psychological springboards in leaping across the logical gap. The resulting inferential process may be one among a wide range of possible types. It may be a generalization from a number of observed facts (*all ravens are black*), the recognition of a pattern (*this looks like the round figure I saw earlier, though with a few kinks at some points*) or a face (*oh, mom, this must be you when you were at school*), or the production of a remarkable

scientific hypothesis (*the missing energy in the decay process of the nuclei must be carried by a mass-less particle that eludes detection*).

An inductive inference differs from a deductive one, depending on whether the conclusion arrived at is necessarily true in the given context of available information or whether the inference expands the information base by adopting a choice from among alternatives, all consistent with the context one started from. Very few inferential acts are strictly deductive in that the starting premises (the overt 'context') are not precisely defined and may have slightly varying alternative interpretations, a large number of premises (possible ones in the given context) having only slight relevance in the context of the problem at hand are ignored, and so on. And, in addition, the *rules* invoked for the sequential processing of information may not be rigorously tested and precisely defined ones.

Inductive inferences are, basically, *informed guesses* and are ubiquitous in individual and social life, as also in specialized scientific inquiry, being almost co-extensive with cognition itself. Most of these inferential acts are carried out, in a great measure in the arena of the cognitive unconscious, in apparent defiance of norms of rationality, with the help of heuristics — half-baked inferences themselves — where one set of heuristics are used to build another, more complex one in a hierarchical process. The heuristics form a part of a vast web of beliefs of various descriptions, where different sets of beliefs have distinct psychological roles to play in the cognitive process. On the whole, the web of beliefs acts as a map for the mind to navigate in a complex and uncertain world, at times sheltering it by making it refuse to accept unsettling facts and, at other times, by spurring it on to take difficult decisions. Beliefs — even relatively remote ones — in conjunction with emotions, affects, and feelings, vastly expand the *context* of an inferential act that often goes unrecognized, especially since the cognitive subject remains mostly unaware of these. It is this expanded context that constrains and guides the inferential act to its desired conclusion. A completed act of inference involves innumerable sequential steps, many of which are branched and tangled, but not many of those follow definite rules that can be assessed in terms of extrinsic normative principles of rationality.

The role of *context* in inference making will be emphasized time and again in this book, where context will be seen to have two aspects to it — the *external* and the *internal* ones. The external context is set by the environmental inputs entering into the inferential process, and is essentially of an inter-subjective nature. For instance, the plaintive cries of a sick child may initiate inferential processes in both of its parents. Here the child's sickness, its past history, its plaintive cries, the household circumstances, the availability or otherwise of a physician on call, all these make up the external context, both the parents being almost equally cognizant of it. On the other hand, the inferential process and the conclusions drawn by them are likely to differ greatly, depending on how they make use of their internal psychological resources. The latter include such diverse things as recollections of past experience, knowledge base, belief system, reasoning ability, emotional makeup and current mood of the two individuals involved. It is to be mentioned, at the same time, that cues from the external context are picked up differently, and to varying degrees, by different individuals, depending ultimately on their internal psychological resources and, moreover, the assimilation of many of these cues occurs tacitly. This is how the external and internal contexts of an inference merge with each other, of which more later.

Within the internal context, factors of a deeply *personal* nature exert their influence upon the inferential process, mostly in the form of *rules* by means of which information is processed within the mind of an individual. Some of the rules used in an inferential process may be universal ones (such as the rules of mathematics), but most are likely to be of a less general nature. Among these latter, there can be found a gradation in respect of the degree to which the rules are specific to an individual. For instance, some of the rules may be in the nature of heuristics or beliefs shared by various different persons (*if an assumption fails to produce a satisfactory result, TRY THE OPPOSITE*) while many others arise in a specifically personal context (*my kid brother is deficient in mental skill, and I must protect him*). What with the huge internal context involved in an inferential process and the non-universal nature of the rules made use of in the process, inference making in general and the making of an inductive inference in particular, is possessed of great complexity and depth, most of which goes unnoticed.

CHAPTER 4. INDUCTIVE INFERENCE

The term 'background knowledge' used by Leighton (in [81]) and other cognitive scientists is an indicator to the internal context of an inferential process.

As I have mentioned, the role of context in the human inferential process, and that of rules constitute recurrent themes in this essay. You will find a summary relating to these two at the end of chapter 7 at [this section](#), as also [this](#).

Chapter 5

The cognitive unconscious

Hidden cognition

You ignore the realm of unconscious meaning at your peril.

Jonathan Lear in [80]

I begin this chapter by quoting the opening paragraph of a relatively recent book on the cognitive unconscious:

“Over the past decade or two, a new picture of unconscious processes has emerged from a variety of disciplines that are broadly part of cognitive science. Unconscious processes seem to be capable of doing many things that were, not so long ago, thought of as requiring mental resources and conscious processes. These range from complex information processing through behavior to goal pursuit and self-regulation.”, [122].

A widely circulated story (relating to the purported advertisement of Coke in a movie theatre) on *subliminal priming*, which later turned out not to be based on facts, was an early indication of the importance of the role of subliminal stimuli (i.e., ones received without awareness on the part of the recipient) in modifying a persons behavior. Investigations have now established that *unconscious pattern recognition* based on priming (an effect where exposure to a stimulus modifies subsequent response to related stimuli)

can indeed occur under certain circumstances ([125], p238-241). In recent years there has taken place a veritable explosion of evidence and data collected in psychological and neuro-physiological investigations on what can be referred to as unconscious cognitive action.

The unconscious has been acknowledged since antiquity in folklore, literature, philosophy, and anecdotal introspective reports by a number of scientists on their own creative processes, as also in early psychological studies and investigations, and was finally made the basis of a *theory* of psychology by Sigmund Freud. That theory was subsequently banished from the mainstream of academic psychology for a variety of reasons but continued as a parallel psychological theory and practice, while academic psychology itself was dominated by the point of view of *behaviorism* which looked upon the unconscious (as well as the conscious) mind as a mere 'construct' on the part of the psychologists that, along with most of the concepts in *folk psychology*, has to be swept aside for the science of psychology to progress and flourish. Behaviorism, however, went through a rapid decline, to be followed by the so-called cognitive revolution where *mental* processes like memory storage and retrieval, and the making of inferences, judgments, and decisions were acknowledged as ones of focal interest, and the interdisciplinary subject of *cognitive science* began to take shape. A major impetus to the 'revolution' came from the field of economics, including investigations on consumer behavior (and studies in the management of business enterprises). The rapidly expanding area of computer science, and the burgeoning field of *artificial intelligence* constituted the other major ingredient of the subject of cognitive science, which provided for the production of *computational models* of human mental processes. This entire complex process of paradigm shift ran parallel to new currents in the field of history and philosophy of science that had an almost unrecognized beginning in Polanyi and that found a definitive expression in Thomas Kuhn.

The *psychology of the unconscious* was brought in within the area of focal interest in cognitive science during the late sixties and early seventies of the last century, and became a trend to reckon with in cognitive psychology during the eighties and early nineties through works of the likes of Arthur S. Reber, Axel Cleeremans, and John F.

Kihlstrom.

Rebers work started from a recognition of insights harvested from Polanyi that led to studies on *implicit learning* and *implicit knowledge* where, significantly, considerations in *evolutionary* biology were of focal relevance ([103], chapter 3), and where the latter underlined the antiquity of the origin of the basic, unconscious, learning mechanisms in the human mind. Cleeremans focused on the neural and computational underpinnings of unconscious cognitive processes [21], while Kihlstrom produced the first influential overall view of the cognitive unconscious [70].

The unconscious as the arena of complex cognitive processes

These early works were then followed by broad currents in the field of psychological investigations in unconscious cognitive processes, and significant areas were opened up one after another, indicative of the ‘primacy of the unconscious’ ([103], p 88). As Kihlstrom summarized the early phase of the work in the area, numerous findings of major significance were already indicative of a broad range of activities of the cognitive unconscious, including the execution of *automatic* processes, based on apparently innate procedural knowledge, subliminal perception (processing of subliminal stimuli), implicit memory, with indications that *implicit learning* was also a distinct possibility, and hypnotic alterations of consciousness.

Subsequent work has made it progressively clear that much of human cognitive functioning that was so long ascribed to conscious, attentive, and deliberative mental activity, was actually carried out in the realm of the unconscious, since the latter was capable of a great deal of *complex* psychological processing. In other words, there has taken place a gradual shift of paradigm in which the ‘old’ cognitive unconscious has been replaced with a ‘new’ cognitive unconscious, where the latter appeared to be capable of even such self-regularity features as intention, motivation, and self-reference [122]. These developments were the result of a multitude of methods and practices that made

possible a great deal of concrete investigations, replacing a more speculative approach of the earlier period.

The unconscious versus the conscious

An issue of fundamental significance that has come up in the wake of this paradigm shift relates to the very concepts of the ‘conscious’ and the ‘unconscious’, since the ongoing work in the field has the tendency of making earlier concepts inadequate. Even the early pioneers were aware that a simple conceptual dichotomy between the conscious and the unconscious was likely to prove misleading ([103], chapter2, chapter 4, [22]). Kihlstrom outlined an early ‘map’ for the ‘substratum’ of the mind by distinguishing between the unconscious, preconscious, and the subconscious [70]. A few authors prefer to use a general term such as ‘nonconscious’ [56], perhaps implying that more specific terms may be needed to describe the structure of what is commonly referred to as the unconscious.

In this book, I will refer frequently to unconscious cognitive processes in the context of inference-making in general, and the making of inductive inferences in particular. As in the case of a number of other psychological concepts, I will entirely confine myself to the discourse of *folk psychology*, in so far as such concepts are *not inconsistent with* findings of rigorous psychological investigations (refer to [114] where Stephen Stich elaborates upon the view that the folk psychology of belief does not deserve a place in cognitive science; contrary views are to be found in [5] and [59]). One has to remember that psychology is itself a subject where, in spite of a vast body of meticulous work, conceptual homogeneity is not of the same order as in physics or chemistry, and various different points of view coexist within the discipline (an analogous, though not comparable, situation in a branch of physics is to be found in the area where quantum theory is sought to be integrated with the theory of gravitation). This leaves room for speculations where, however, there has to be *consistency* with the massive body of findings of an evidential nature. Reber famously used the phrase ‘sensible speculation’ ([103], chapter 3) in the context of discourse on implicit learning in relation to evo-

lutionary theory, because he based all his speculative observations on findings from the current body of investigations over a wide area. This book of mine contains, at places, speculations of a much more dubious nature since I have had little direct acquaintance with journal articles, which pour out by hundreds every day, in the fields of cognitive psychology and philosophy of science, and I base most of my observations on monographs in the fields, taking care not to be inconsistent with these. I will say no more since I know that disclaimers are odious.

Unconscious cognition: the role of emotions

My aim in the present chapter has been to highlight that, according to current views on the subject, the unconscious mind is capable of a great deal of complex cognitive activities, including ones where ‘cold’ cognition is intimately blended with emotion-based factors (giving rise to what is, at times, referred to as ‘hot thought’ [118]). There exists a large body of literature on the possible role of emotions in cognition, based on evidence of substantial value, and part of that literature pertains to emotions and affects in the workings of the cognitive unconscious (see, for instance, [35], [34]). Much of this book is focused on the possible role of beliefs, emotions and affects, operating within the realm of the unconscious, in inference-making in general, and in inductive inference in particular (see chapter 6).

The unconscious detection of similarity

One particular focus of interest will relate to *hypothesis*-building, conceptual transformations, and creative episodes in science which, within the framework presented in this book, constitute a very special kind of inductive inference (see chapter 8). The emergence of hypotheses depends crucially on an important organizing factor in the establishment of *correlations* among apparently remote ideas, and in the building of novel conceptual structures out of these, principally by means of the detection of *similarity* at the unconscious level.

Similarity is a notion of great relevance in cognition (refer, once again, to chapter 8).

The detection of similarity between perceptual inputs and the formation of appropriate inferences on the basis of such similarity is a basic mechanism in survival and natural selection. The mechanism underlying the capacity of such similarity detection is, perhaps, of an elementary nature based on associative covariation detection between the perceptual inputs. However, it is likely that the cognitive mechanism has the ability to induce rules of inference in a hierarchical manner, where rules at a lower level of complexity are made use of to induce those at a higher level, and where the associative detection of covariation among a series of successive perceptual inputs is likely to be one among the rules of inference at the lowest level. Based on the elementary capacity of detection of similarity of a coarse nature between objects in the environment, more complex and fine-tuned rules of similarity detection, first among *objects*, and then among *concepts*, are likely to be induced by the cognitive unconscious, resulting in a heightened capacity for the formation of novel conceptual structures – the *hypotheses* that germinate into *theories* in the ongoing process of scientific exploration (for background, refer to [62]).

The ‘rationality’ of the unconscious

The cognitive unconscious is commonly contrasted with the *conscious* mind, and understandably so. While spontaneous information processing activity within the unconscious mind of an individual continues without awareness and without overt intention on her part, conscious processes are supposed to be executed deliberately and intentionally, within the horizon of awareness of hers. Indeed the *two-process* theory of human cognition and rationality (see chapter 7 below) is based on the assumption of two *types* of processes being involved, where the two types correspond to contrary and complementary features of inferential activity reminiscent of those commonly ascribed to the unconscious and the conscious levels of the mind. However, this approach, along with the concomitant use of terms like ‘levels’ or ‘strata’ (the unconscious being a ‘substratum’ of the conscious), constitutes no more than a convenient description of processes that is, in all likelihood, a somewhat simplistic one ([103], chapter2, chapter 4, [22]), though convenient and useful too. One has to mention at the same time that

the proponents of the two-process theory do not always associate the two processes with unconscious and conscious cognitive activities, confining themselves more to the consideration of *computational* features of the two.

In this context, a number of relatively recent findings shed a new light on the so-called conscious-unconscious divide that may prove to have far-reaching implications. Briefly stated, the conceptual attributes of consciousness may need substantial revision. As the body of findings continued to grow on the capabilities of the unconscious mind for various types of complex cognitive processing, it appeared that what goes by the name of consciousness has a correspondingly small role to play in cognition as a whole, namely one confined to control and regulation of the cognitive process. For instance, within the framework of the two-process approach, it may conceivably play the role of a ‘decision-making entity’ performing the job of planning and review, such a decision-making job being compatible with the commonly held view that the conscious mind is capable of setting goals, and of working on the basis of an *intentionality*.

One central issue of the conscious-unconscious divide relates to the computational question of the role of a working memory. It has so long been a basic assumption that the conscious mind differs fundamentally from the unconscious from a computational point of view, where the former involves a working memory that sets goals within a serial information processing framework, based on more or less well defined rules or principles, and that such processing requires attention, as also intention ([68], p65- 66, [125], chapter 2). By contrast, the processes within the cognitive unconscious were assumed to be based on distributed transitions of an associative nature, without the involvement of a working memory. It now appears that, from the point of view of functionality, the unconscious mind possesses the ability to set goals and to hold information, resembling the role played by working memory [122], [55], [115]. In particular, the *parallel distributed processing* (PDP) model (or the *connectionist* model), based on the computational abilities of neural networks offers a great many features, including those relating to a working memory, even without the necessity of assuming the existence of a separate specialized memory unit. Though the PDP Model ([21], [10], [11]) is basically a computational one and does not enjoy firm backing within the tradition of psychological

research, it can nevertheless be used as a good metaphor, providing pointers to a host of possible functional features of the human cognitive process, including the ‘rationality of the unconscious’ (see, for instance, [91], [44]).

CAUTION! The note of warning that I want to record here is that the ideas on unconscious cognitive activity are not part of a solidly based theory, and much of these are conjectures and plausible speculations though, at the same time, parts are indeed based on conscientious experimental findings. There is an ongoing controversy on how much of the theory of unconscious thought is admissible on evidence and how much is not. The framework for a theory of inductive inference that I seek to outline in this book does contain conjectures and speculations that are in the nature of interpretations on my part, not inconsistent with more disciplined theoretical and experimental findings. I have stuck my neck out, without regard to whether the guillotine of the science of cognitive psychology will descend on it and chop it off.

An instance of the ongoing controversy on unconscious thought theory referred to above is to be found in [65].

Neuropsychological studies over the last few decades have revealed, and continue to reveal, a great many facets of neuronal activity bearing upon complex psychological processes that appear to be predominantly unconscious. These relate to affects, emotions, and feelings on the one hand, and cognitive processes like belief-formation, decision-making and planning on the other. A number of neural aggregates of early evolutionary origin, responsible (in conjunction with a number of chemical secretions in the brain) for the generation of *pleasure*, *reward*, and *aversion* have been found to play a crucial role in these complex psychological processes (see, among an extensive literature, [89], [8], [31]). Finally, consciousness appears as an epiphenomenon of sorts, resulting in capacities such as the ones of awareness, intention, and introspection, where large scale synchronized neuronal activity along with the mediating action of a set of *neurotransmitters* are of central relevance [3], [4].

Unconscious cognition: how ‘hot’ is it?

I close this chapter with a few words on the possible relation between the *cognitive* unconscious and the *Freudian*, or *psychoanalytic*, unconscious. In this book, I will not make overt reference to the latter, in deference to the major part of literature in cognitive science which has a tendency to maintain a distance with the Freudian framework, perhaps because the latter is still felt to be too ‘hot’ — too infected with concepts of a primitive nature and with the idea of a primeval mind — to be accommodated within a respectable cognitive theory. The early phase of development of cognitive science was indeed one where it had a ‘cold’ and ‘respectable’ face. However, developments of a more recent vintage have made it imperative to take seriously to the idea that human cognition is crucially dependent on affects and emotions, and is therefore not as ‘cold’ as one would like it to be ([118]; [68], chapters 5,6). Even when one discounts a considerable part of the Freudian picture of the unconscious — the one that has turned out not to be solidly grounded — the psychoanalytic *point of view* may still prove to be useful in the working out of an integrated view of the cognitive unconscious, where a wide range of human emotions, including motives, desires and drives of various descriptions (where, moreover, ‘pleasure’ and ‘reward’ centers in the brain play their role, aided by a number of chemical secretions) are involved in a cognitive process whose ‘rationality’ will have to be interpreted anew.

Chapter 6

The process of inference: beliefs and emotions

The belief-laden quest

Remember that all perceptions, judgments, and beliefs are inferences and not direct readouts of reality.

Richard Nisbett in [91]

Beliefs and belief systems

Inductive inference is, essentially, the formation of *belief*. Myriads of external and internal stimuli impinge upon one's perceptual mind every moment, of which some are selected out, in keeping with the *context*, and processed in association with the latter, to be eventually transformed into beliefs and to be incorporated into the belief *system* of the person concerned.

Incidentally, the issue of context is of great relevance in inference, decision-making (which I will count essentially as a form of inference), and the formation of belief, and has already made its appearance in various different places (various different *contexts*,

I am tempted to say) in this book. Broadly speaking, context can be of the *external* or *internal* type, analogous to stimuli received by the mind. When I am looking at a beautiful portrait, the canvass, the gallery, the people around me, the artist, the occasion of my visiting the gallery and viewing the portrait, and a thousand other things make up the external context of my act of viewing and appreciating. On the other hand, my past recollection of portraits of a similar kind, my admiration for the artist, the faint similarity of the face drawn in the portrait with my mother's, my slight feeling of discomfort at the color contrast in one part of the portrait, my elation at the sense of vitality and jubilation expressed through the portrait, my current mood of anxiety and depression, and a million other things in my mind and my emotions constitute the internal context.

Many of these things, in some form or other, will make the content of my future memory of the portrait and of my act of viewing it, and will be blended into my subsequent recollection of that memory. And, interestingly, that recollection will again depend on the external and internal context of the act of recollection. In addition, my viewing of the portrait, may lead to the formation of a belief to the effect, for instance, that age is finally catching up with the artist I admire as one of the foremost among the current generation. Here again, the belief is formed in a process of great complexity occurring amidst the external and internal context I mention above. The complex issue of the context in an inferential act has been and will be a recurring concern of ours in this book.

To come back to the question of beliefs and belief systems, beliefs constitute a most elusive entity in cognitive science [24], even as our entire mental life is permeated with beliefs that relate to almost all aspects of our thought process. Beliefs can be of various kinds and various shades, ranging from mild predisposition to virulent dogma, from unconscious mental states to consciously held points of view, and from justified and true ones to overtly inconsistent stands on issues of various degrees of relevance in our life. Beliefs are arrived at in response to situations that a person or a group of persons faces, and are used as *guide maps* in subsequent response of that person or group to situations and scenarios that come up from time to time. Our belief system is

our map to navigate in this largely unknown world [2] and as such, the beliefs have to have an optimum degree of permanence. If all our beliefs get revised under the slightest challenge from the world around us, then the belief system will be a very poor map indeed, and we would face disaster in trying to navigate with its help. If, on the other hand, all our beliefs were rigid and resistant to revision, then again these would be of no use because these would then prove unequal to the task of making us aware of the complex and changing realities of the world. In fact, some of our beliefs are extremely resistant to revision while some others are easily revised under the impact of facts of the world and of rules of tested efficacy learnt from past experience (to be precise, these rules are also mostly in the nature of beliefs) and eventually gain the status of *knowledge*.

In the present section we will mostly be concerned with beliefs and belief systems of individuals. Beliefs of an inter-subjective nature are also of great relevance in the collective life of communities and entire societies, where these form integral parts of culture, religion, folklore, rituals, and taboo. Communities of scientists also carry their beliefs in various forms. This raises the question as to how much of scientific theories are in the nature of social constructs. Or, to take up a much broader issue, how much of science in our times is socially determined? We will look at some of these issues in subsequent sections in this book.

Belief, knowledge, and inference

It is of some interest to examine the relation of beliefs to conclusions drawn in acts of inference and, also, to items of knowledge. Of these, let us have a brief look at the latter relation first. It is sometimes said that knowledge is justified true belief, and I will not join issue with this point of view, though the questions of justification and truth are contentious ones. As I have mentioned, knowledge emerges out of an evolutionary process from belief in the course of revision of the latter under the impact of reality. However, we have already had a glimpse into the *layered* structure of reality, which makes it a fathomless entity, and there can, truly speaking, never be anything

like ultimate knowledge (more of this later, because this is a central issue in scientific realism). With this little rider on the concept of knowledge, I should like to draw your attention to another, more subtle, aspect of the relation between belief and knowledge, and that happens to be, in a sense, a *psychological* one. Beliefs mostly remain inactive or latent in our mind, and become active when we ourselves need to act in some way or other, say, under the impulse of fear, or craving for recognition, or of any other cause. The course we adopt in embarking on an action (which may, under special circumstances, even be an act of inference) is shaped by a set of beliefs becoming active and guiding us through in that course. Beliefs propel us into action (*the tyrant is the root of all evil*) and guide us through (*smash the palace*), but knowledge has no such role. In the course of evolution of belief into knowledge the latter, in a manner of speaking, gets sterilized and cannot infect the mind with germs that make it febrile and pro-active. We *make use of* knowledge in the course of action, but we are not *spurred on* by it. That role is vested in belief. Looked at another way, beliefs are, generally speaking, associated with *emotions* (more of this below) and, in the process of evolving into knowledge, get divested of these emotions.

Indeed, the web of beliefs in a persons mind, along with the associated network of emotions, can be said to identify the *self* of that person. In contrast, knowledge is like an encyclopedia that she carries along with her and makes use of as the necessity arises. Our closest and most intimate beliefs are tied to our innermost selves by means of emotions. These are as completely resistant to revision as our own selves are to change.

An inferential act is an act of a special kind where we are activated toward a goal or a purpose such as the purpose of solving a puzzle or a problem (*how to calculate the energy of binding of this crystal?*) where, generally speaking, the purpose is shaped by some event or situation that acts as the trigger for the action (*have to have a Ph.D.*). The trigger makes us aware of the necessity to solve a problem or to achieve a goal or purpose, and this sets in motion the mental process of making an inference. The process needs a number of inputs to proceed to its desired end. These are commonly

referred to as the *premises*, while the end point of the process often appears in the form of a *conclusion*. For instance,

I am feeling indisposed / physicians need be consulted when one feels indisposed / there is a physician nearby whose telephone number I have in my diary / the physician is to be called over the phone.

This is a particularly simple form of an inference, where the inputs are in the form of propositions, and there is a well defined line of reasoning that leads me from the premises to the conclusion which is, strictly speaking, one in the form of an implied belief (*he will cure me* — a belief that activates me to make a phone call) but is, within the given context, a justified one.

Not all inferences, however, are of this simple type indeed, very few are. Instead of propositional premises, the inferential process starts with external and internal mental inputs, possibly in the form of some kind of *representations* in the mind, and these are then *processed* (i.e., taken through a succession of transformations) so as to lead to the ‘conclusion’ which is again, more often than not, in the form of a mental representation or mental state rather than of a proposition. This mental state is, generally speaking, in the nature of a belief that, at times, sets one in the course of some action. This said, I must add that not all beliefs are formed by way of inferences aimed at achieving some goal or other. Beliefs are also formed *passively* as an individual (I have already mentioned that the beliefs of groups or communities will not be explicitly referred to for now) passes through experiences or situations. On finding an ailing old man lying by the roadside, I walk by him and find that others are doing the same, and a belief is formed, to the effect that *this world is a heartless one*. The only difference between the formation of this belief and a process of inference as sketched above is that here the belief is formed without an apparent purpose. While I highlight this distinction here, it is an arguable one since some underlying purpose or relevance may still be there, perhaps at an unconscious level, without being apparent (incidentally, *heuristics* are also little beliefs that are often formed spontaneously, without apparent purpose, by way of association of environmental stimuli or of ideas, mostly in the unconscious mind, to

be made use of in future inferential acts). One other point of distinction characterizing the instance under consideration is that the belief is formed as a latent one, without setting me on to a course of action, though it may be instrumental in doing so at some future point of time. On the other hand, a belief and an inference are both produced by a processing of information in the mind. At the risk of some controversy, I may say that the end product of an inference is a belief, while a belief is not necessarily arrived at by way of an overt inferential act.

In particular, beliefs and belief systems are intimately associated with acts of *inductive* inference. In an inductive inference, the mind works in a halting, groping way, without being guided in a clear-cut manner by *rules*, where the rules may be of a precise, logical nature or may be more diffuse ones, arrived at on the basis of prior experience and knowledge. A rule-driven process will be referred to as *reasoning* and is one where there is relatively less scope for guesswork that takes one *beyond* what follows from the starting premises by the application of the rules. In reality, no inferential process is completely rule-driven (we have seen how even *mathematical* reasoning, or reasoning acts of a deductive nature, involve informed guesswork typical of inductive inference) or is totally unassisted by rules. In other words, inferences are part deductive and part inductive, which is why the end product of an inference is, generally speaking, a belief — one that is fallible (or, is *non-monotonic*, if I may use a jargon), is arrived at by a mental process in the nature of a guesswork, and makes use of the context in a manner that is largely indeterminate. And, the other area of relevance of beliefs in inductive inference, relates to this question of context.

As I have tried to indicate in earlier paragraphs in this book, the context of an inference is a highly complex matrix made up of diverse factors. We will be concerned here with the *internal* context of an inference, which assumes relevance when the mind does not have clear-cut rules to make use of while being engaged in an inferential act. It then falls back upon the hidden context, the resources available within its own tucked-away treasure house, howsoever remote these might apparently be. What are the resources that can possibly be of relevance? Here the list is virtually inexhaustible and it is here that we face a very deep issue.

While walking on a broken road, you are likely to come across a breach on which no plank or temporary bridge has been set for your benefit. What would you do to navigate the breach? You will jump across it or, if the breach is too wide for that, you get hold of a stick or a pole or of a branch of tree, whatever comes in handy, and you leap across, using it as a prop for propelling yourself. Likewise in inference. The mind gropes for a prop and makes use of whatever comes in handy — most likely, some belief or other, since beliefs are always there in abundance, associated with one another in a vast web, much like a social network where you just have to cry out for help and somebody or other, though perhaps of dubious credentials, comes forward.

Beliefs and emotions

What is more, beliefs are often associated with *emotions*. It is precisely this association that makes a belief a device for spurring an individual into action — whether an act of aggression or one of submission, or one of a different kind, i.e., more generally, into taking a decision and acting upon it. And it is this association that makes emotions play a vital role, both in a positive and in a negative sense, in cognition. The stronger the emotional association, the more resistant to revision is the belief while, on the other hand, this very association helps the owner of the belief in her cognitive act of adopting decisions, possibly without conscious awareness, in the course of making an inference.

Aspects of the complex relation between emotions and beliefs have been discussed in [41]; [39] speaks of beliefs as being located ‘at the interface of affective and cognitive processes’.

An inferential act in real life differs from a conventional computer program in that decisions are to be adopted at several points in the course of the inference, based on judgments, for which sufficiently compelling inputs are not available. A judgment is, fundamentally, an extra-logical act because it involves the comparison of disparate and, possibly, conflicting items that cannot be reduced to a common denominator. I have a limited amount of money at my disposal as I enter a departmental store and debate

whether to purchase a cosmetic item for my wife or a toy for my child. The two options are disparate, and my judgment and consequent decision can never be *logically sound* — indeed it has to be dictated by incidental factors such as, for instance, how much more I value the smiling face of my child than that of my wife. In other words, it is my emotional make-up that plays proxy to my reason.

So, here is a likely scenario. As I proceed in the course of making an inference, there occurs a processing of information that builds around a sequential application of rules, where the rules can be of various types and descriptions — mostly heuristic ones that have been formed by past experience in inference-making or, in other words, in *interpreting* the world we find ourselves in. However, the sequence is not a neat, linear one but is tangled in a complex way, where there are branchings, backtrackings, and inter-connections between parallel processing sequences. Within this complex and tangled information processing scenario, I face from time to time, an impasse, where my repertoire of rules and reason prove inadequate in the task of helping me along in making a judgment and adopting a decision. I inwardly grope for some clue, something in the storehouse of my mind that I can use as a support in my mental leap through the impasse, and I try to choose from a set of heuristics of, perhaps, dubious value for the task at hand. I make a number of trials with these and in each such trial I seek the help of background beliefs for guidance, and a wide web of beliefs presents itself. As I am about to apply one or more heuristic rules, one or more of these beliefs come forward to either approve or disapprove of my impending act, by invoking an *insidious emotional reaction* — related to some kind of pre-conscious affect. This reaction or affect acts as a signal, of either a positive or a negative nature, and helps me navigate through the impasse by either proceeding forward or backtracking and making a renewed effort. Neither the belief nor the affect is an infallible guide in my decision-making and I may eventually err in my inference. But inferences are never foolproof. What is important is to be able to *arrive* at an inference because, once arrived at, the inference can be tried, tested and improved upon or, in the worst-case scenario, may be rejected out of hand, to be replaced with a new inference — a new belief if you like. Shorn of the *evaluative action* of our beliefs and emotions as the cognitive mind invokes heuristics in bridging a logical

gap, cognition gets paralyzed in its inference-making efforts.

Significantly, the fact that our inferences are guided by heuristics, beliefs, and emotions of dubious authenticity, does not make the inferences entirely devoid of credibility and validity, since all these heuristics, beliefs, and emotions are products of past experience. While some of these are well founded and some others are distorted interpretations of reality, their overall effect is to make us capable of an inference that is, frequently, not far from the truth. Inferences never cease. Once formed, an inference is not consigned to the cold storage. It continues to be tested and modified, if not within the visible space of conscious awareness, then most certainly in the invisible world of the unconscious. It is applied upon the world out there or subjected to consistency checks with our knowledge base and with the existing web of beliefs. In other words, our knowledge base, our beliefs, our inferences, and our emotional set-up are in a constant state of dynamic interaction — interaction among themselves and interaction with the world around us. The upshot is an enrichment of our knowledge base, an enrichment of the vast storehouse of heuristics, and a rebuilding of the great web of beliefs, with some of the beliefs acquiring the credibility of knowledge, some remaining in a state of suspension as ones of uncertain credentials, and some transformed into weird dogma. Human cognition is an ever-active process — and is complex indeed.

Inference is not and, indeed, *cannot be* a neat linear sequence of information processing because such a sequence is possible only when the cognitive system is certain that it will lead to the desired conclusion. Real-life inferences are never made that way, except in the arithmetic class-room. Instead, at every step of reasoning and inference-making, the cognitive mind fans out in the form of *parallel* processing of information — much as the experienced general sends out groups of detachments so as to locate and pin down the enemy. In pursuing all these parallel branches of information processing, each by means of an application of clusters of heuristics, the cognitive mind continually carries out consistency checks against bits of stored knowledge and against heuristics, beliefs, *and* reason, thereby approving of or discarding one or more of these branches with a view to achieving some desired goal or purpose. In other words, there operates an ‘internal censorship’ in every inferential process, as Medawar tells us in the following

passage:

“In real life, of course, just as the crudest inductive observations will always be limited by some unspoken criterion of relevance, so also the hypotheses that enter our minds will as a rule be plausible and not, as in theory they could be, idiotic. But this implies the existence of some internal censorship which restricts hypotheses to those that are not absurd, and the internal circuitry of this process is quite unknown. The critical process in scientific reasoning is not therefore wholly logical in character, though it can be made to appear so when we look back upon a completed episode of thought.”,
([88], p 53).

The internal censorship that Medawar speaks of is not really ‘quite unknown’, since much is now known about the ‘reward’ and ‘value’ systems, based on neural aggregates, in the brain [31], [8].

In the process of parallel exploration and concomitant evaluation of inferential possibilities, branches get connected and tangled, with the tangle eventually assuming the form of an inference — an inference that is seldom a sharply defined proposition since it is mostly in the nature of a hazy concept associated with other concepts and beliefs in the mind, though sufficiently defined so as to be made use of in acting upon the external or the internal world.

Imagine yourself engaged in some complex mathematical derivation, where you apply the *universal* rules of mathematics and logic but have to decide from time to time as to which of several alternative courses of application of the various rules known to you are to be applied so as to successfully complete the derivation. Strictly speaking, you cannot examine all possible courses of application of these rules, if only because that would entail a prohibitively complex enumeration of the possibilities, and you allow your mathematical *instinct* to come to your help where the instinctual help comes in the form of a large number of heuristics, mostly bits and pieces of mathematical reasoning themselves, floating about in your unconscious mind. While you are proceeding with

the derivation, the cognitive unconscious is not sitting idle, since it is engaged in small excursions and forays in parallel to your main line of derivation so as to tell you which course to adopt whenever a decision is to be taken. But now imagine that you have got yourself blocked at an impasse where your storehouse of mathematical knowledge and your repertoire of heuristics of a mathematical nature fail to bail you out.

If you are circumspect by nature, and rely too much on secure mathematical reasoning, you will stare and stare at your worksheet, unable to proceed further. But your unconscious mind is not circumspect, and it does not sit back staring. It turns and looks at beliefs. In some distant past, you came across a derivation made by a peer in a remotely related area where the problem at hand, when looked at from the point of view of a *higher dimensional space*, got solved in an unexpected manner. You now make use of a heuristic that launches you on to a higher dimensional space, and you get the nod of approval from your belief that what the peer did would bring dividend to you as well. Luckily, your hunch does pay dividends, but the more important point is, it got you released from your mental block so that even if the hunch were to prove unproductive, you would backtrack and then make good use of some other hunch. This is a hypothetical instance that I include here to bring home the point that the reasoning we employ in an inference does not always involve rules of the same degree of rigor — while some of it is made up of universal rules of mathematics or logic, some may be of an ‘inferior’ quality, and may even be in the nature of loosely woven beliefs, but it is the web of belief that keeps the inferential process going when reasoning of more reputable pedigree fails, and, through excursions, forays, and parallel processing, helps the process to reach a destination, somewhat like a drunken person finding home (*perhaps* his own) at dead of night. The cognitive mind is opportunistic by nature and does not hesitate to get drunk with disreputable stuff.

As all of you must be aware of from your own experience, the formation of a belief is often associated with an emotion specific to the situation in which it originates and, so is its activation at a subsequent point of time. And, emotion and affect are not merely psychological entities, since they involve physiological and somatic factors as well. One thus has a highly complex system to look for when the question of the *context* of an

inferential process comes up.

Emotions and affects are all-pervasive in the mental life of a person. These find expression as responses of the entire mental and bodily system to situations of diverse types — responses at times intense and at times barely sensed. Likewise, a person's web of belief is based on neural co-ordination on a large scale is not localized in this or that neural aggregate (on the other hand, beliefs of various *kinds*, such as those evoking responses of assent and dissent, are possibly associated with corresponding neural aggregates; see [107] for a number of interesting observations) — it permeates his entire psychological being, weaving through the entire body of his conscious and unconscious thoughts. And this entire system of enormous complexity is potentially the 'context' of an inferential act. It is from here that the mind draws its ingredients and springboards for negotiating and vaulting over the logical chasms it encounters in the course of an inferential act.

Thus, in other words, an act of inference has strange and hidden aspects to it that cannot be captured in commonly invoked norms of rationality. The deeper and the more complex the substratum of the cognitive mind an inferential process gets drawn into, the more deviant it appears from reconstructions based on *extrinsic* rules by which the process is appraised and evaluated.

The role of emotions in cognition is a topic under extensive investigation, and has now generated a quite voluminous literature, inclusive of [1], [63], [100], [79], and [47] (in addition to references already mentioned).

Heuristics in the inferential process: the dual role of beliefs

Speaking of the 'hidden logic' of an inferential process, one comes face to face with the role of heuristics in inference making. Heuristics, as I have mentioned, are rules of thumb that the mind makes use of in making judgments and decisions and in forming

inferences, in which process it is often rewarded with success, but failure is not ruled out either. Heuristics are little beliefs of ours that keep on being formed, at times, without conscious awareness on our part, though we may subsequently become aware of some of those and even make use of those (*in the endgame, play the king*) while solving problems or making inferences. The scope of applicability of a heuristic may be limited to only some specific problem situation or may be a comparatively extended one spanning a whole class of problem situations. And, heuristics often work in tandem with one another, and in clusters, helping and prodding us along in the course of an inference. Heuristics are, in a manner of speaking, hidden *rules* operative within the mind, though mostly of an *intrinsic* nature, many of those specific to an individual — specific to her developmental history, to her response to innumerable real life situations in the past, and to her repertoire of experience tied to emotions and feelings.

From a computational point of view (this is the view adopted in the theory and practice of *artificial intelligence* that tries to capture various aspects of human thought, learning, reasoning, and intelligence, thereby shedding light on human cognition at large), heuristics are the counterparts of *rules* and *rule clusters* made use of in computer realizations of inductive inference as described, for instance, in [61], where a rule typically operates upon mental *representations* of the internal and the external world, and can typically be described as an *if-then* operation upon such representations. The rules operate hierarchically, where an intermediate inference, arrived at on the basis of one set of rules, itself acts a rule for the next stage of inference, thereby realizing a complex and tangled sequence of ‘information processing’ whose ‘code’ is formed and lodged within inaccessible recesses of the mind.

Heuristics are in the nature of *clues* to solutions of problems — typically, clues picked up from the context within which a problem is defined, where both internal and external contexts are made use of by the reasoning mind, mostly in a tacit manner. As I have already mentioned, an inferential process involves the operation of rules more or less *independent* of the context specific to the reasoning individual, *as also* clues generated from the context that is specific to the individual. A heuristic is a rule, mostly of a tentative nature, that can be of either of the two types. Thus, there are heuristics made

use of by most individuals in some given situation (*dont ever argue with the boss*), while there are others that do not transcend the individual and are not shared by others, the latter being, mostly, heuristics lodged in the unconscious mind.

In a manner of speaking, heuristics can be described as beliefs that are made use of as ingredients in inferential processes and, commonly, are ones that are subjected to consistency checks to a greater extent as compared with many other beliefs that are resistant to revision. It is this fluidity of the heuristics that makes them useful in the generation of valid inferences. These are clues often produced ‘on the fly’ as the process of inference progresses, and some of these are also discarded in a similar manner as they fail to guide the reasoning mind properly toward an acceptable solution to the problem.

But how can it be ascertained whether an inference is progressing toward an acceptable solution at an intermediate stage of the process even before that solution is arrived at? Again adopting a computational point of view, one can conjecture that a process of inference is not just one single sequence of information processing, but a large number of sequences proceeding in parallel, as the mind tries out a number of alternative inferential pathways *simultaneously*, with diverse sets of clues of objective and subjective natures (the former being clues shared by different individuals in a given situation, and the latter specific to the internal context of an individual). Most of these parallel paths of inference are discarded as it transpires, through consistency checks of various kinds (where such consistency checks often involve *counterfactual* thinking [27]), that these are not leading to any satisfactory solution to the problem at hand (*this approach is getting too complicated — leave it and try something simpler*), while the remaining ones are pursued till one or a few of these seem to engender a solution, or a number of alternative candidates for a solution, appropriate to the need of the hour. This must be how the mind ‘scents’ the right path to the solution of a problem at hand. At every stage of an inferential process, and at each step of all the parallel inferential paths adopted in the process, the reasoning mind has to *judge* and to *decide* between *alternatives* in the context of a set of consistency conditions that an acceptable solution to the problem should meet with. However, these conditions are not always determined in a clear-cut

manner in terms of subject-independent ones.

And *here*, as I have briefly explained above, is to be found a second, distinct, role of beliefs in the inferential process, a role that is intimately associated with emotions and affects. Thus, at times, as the mind prepares to take a logical leap, some belief lodged in it comes forward as an arbiter or evaluator as to whether that course is an 'acceptable' one. If the course of inference is at variance with the belief then a warning bell is sounded whereby the mind shrinks from pursuing the course in question by experiencing a negative affect. What is of relevance here is that the clamor of the metaphorical warning bell comes into being as a consequence of *emotions* associated with the belief that comes forward for the appraisal as to whether the inference is proceeding in the right direction. At times, the belief in question appears to be in consonance with the course of the inferential process, the direction in which the mind is ready to leap for making possible the onward progress of the process, when a pleasurable affect waves the process on. This entire thing is, of course, contrary to what goes by the name of logic since, in terms of logic, the belief that appears in the role of arbiter or evaluator may not even be of direct relevance to the problem at hand. Suppose, in the course of deciding upon the number of invitees in a party, the number twelve comes up as a possibility, but the person in charge making the decision has an abhorrence for that number — an abhorrence of an entirely personal nature (perhaps his son died on the twelfth of a certain calendar month) that has got nothing to do logically with the considerations necessary for making the party a success. Still, the prospective host recoils from the number and looks for ways to extend or to shorten the list of invitees (the case would have been different if the number in question were thirteen instead of twelve because then the dread of that number, though still an 'illogical' one, would have been in the nature of a 'rule' shared by many individuals, and would resemble a 'logic', based on a rule).

A traumatic experience like the one mentioned here often generates a set of beliefs, associated with emotions of an elemental nature, that have an overriding influence in setting the course of an inference, to which these beliefs have even a remote relevance.

Though, in the present instance of application of a rule (either an inter-subjective or a person-specific one — these terms are further explained below), the drawing up of a list of invitees does not have to satisfy the needs of some objective *truth*, still it illustrates how ‘rules’ of diverse kinds make inroads in our inferential process.

More generally, inferences are made necessary so as to guess at ‘truths’ of the world or to decide upon a course of action in order that some goal or purpose is achieved. In order that the desired goal is actually arrived at, inferences must have a commitment to the external reality — to correctly grasp some aspect of that reality. Rules of various kinds are routinely invoked in the cognitive mind to guess at such truths.

Briefly stated, as a juncture is reached in a process of inference where the reasoning mind is poised for a decision as to the possible efficacy of a heuristic in bridging a logical gap, the likely consequence following from the adoption of that heuristic (or of a cluster of heuristics) is judged against a belief (or, again, against a set of beliefs), where the result of the judgment appears as an affect or a feeling of either a positive or a negative kind, in virtue of some associated emotion. An emotion, generally speaking, is an *amplifying* mechanism designed to make a subject sensitive to a situation, whereby she can either take evasive action or one of engaging into the situation with a positive frame of mind. It is this *amplifying and evaluative* mechanism involving beliefs tagged with emotions that is possibly of crucial relevance in helping the reasoning mind navigate logical gaps in the course of an inferential process.

In summary, beliefs play a dual role in inferential processes in general, and in inductive inferences in particular. First, they are used in the form of heuristics in the step-wise progress of an inferential process, much in the same way as logical rules are made use of in a mathematical derivation, with the great difference that these are themselves in the nature of products of guesswork, and are often of dubious efficacy and, above all, may be highly subjective in their origin. The *other* role of beliefs is that of acting as evaluators and arbiters (‘internal censorship’) in the inferential process when judgments and decisions (involving a choice between disparate and conflicting alternatives) are to

be made in navigating across gaps in which the use of heuristics, just by themselves, is not sufficient for the success of the inferential act. Here a belief, in conjunction with the associated emotions, makes the reasoning mind alert as to whether the course it is going to adopt is about to bring it nearer to the desired goal or to take it further away. This role of beliefs is a curious one since the belief in question may not be of direct relevance to the goal that the inferential process is to culminate in, but is still made use of as part of an overall strategy in a broader context.

For instance, here is a hypothetical situation that I include more as a metaphor than as an illustration. Suppose that a mother is playing a game with her child, asking her to search out an object which she, the mother, has hidden somewhere around. At one stage in her exploring spree, the child comes to a momentary halt, unable to make up her mind as to whether to enter into the kitchen or into the drawing room, when her mother issues a shrill command that she is not to go into the kitchen. Here, it is quite logical for the child to assume that the kitchen may be a possible hiding place, but the mother's admonition has a different relevance, in a context broader than the issue at hand.

Deterministic but unpredictable

The mind is a system of great complexity. One can look at an inferential process, or any process in which the mind participates, as a *trajectory in a state space* of an enormous number of dimensions (I will come back to this idea of exploration of a conceptual space at greater length in chapter 9). Arguably, the trajectory remains confined to a subspace (in a loose sense) of a smaller number of dimensions in any specific process of reasoning and inference, but that process may still be *dynamically complex* in the sense of being *unpredictable* while still being *deterministic*, where the latter term is used in a broad sense, meaning a process that is predominantly driven causally from one step to another by rules, though the latter may be rules that have a specifically *subjective* component, and ones that are fallible.

Heuristics as inter-subjective and person-specific rules

Two terms that assume relevance in this context, and have featured in earlier paragraphs, are: *inter-subjective* and *person-specific*. Thus, rules of an inter-subjective nature are ones that are shared by individuals going through an inferential process defined by a given *external* context (external inputs, triggers, and goals). Referring back to an earlier example (a completely hypothetical one), if a number of persons are asked to draw up plans, each in an individual capacity, for arranging a party, several of them may avoid the thirteenth day of the month as the projected day on which the party is to be held, though that same day may come up as a likely possibility as a result of a few other considerations (it may, for instance, be a national holiday). Here the abhorrence for the number thirteen is a belief of an inter-subjective nature (which, however, may still not be a *universal* one; rules may differ in their degree of universality). On the other hand, one particular individual may have an abhorrence for the number twelve because of a tragedy of a deeply personal nature that is associated with the recollection of the twelfth day of a month, and he may even prefer the thirteenth day over the twelfth for that reason. What is important to note here is that both inter-subjective and person-specific rules are, under most circumstance, *fallible* in nature, being products of guesswork and of prior beliefs which, however, does not prevent the reasoning subject from making use of these in his reasoning and inference-making.

The question as to their efficacy in arriving at ‘truths’ is a deep one; even though fallible, the rules are linked to past experience and, when used in conjunction with rules of greater efficacy and the ‘internal censorship’ in our mind, they often guide the cognitive process to a point not very far from a successful conclusion.

The issue of ‘rules’ in inferential processes has been addressed by Reber in [103] in a manner at once subtle and nuanced.

According to the framework outlined above, heuristics play the role of inter-subjective and person-specific rules made use of by the reasoning mind in the course of an inferential process, while, generally speaking, a number of inter-subjective and person-specific

beliefs are also involved in the process for the purpose of judgment and evaluation at various stages of the process, resulting in the decision as to whether the use of this or that heuristic in the bridging of a logical gap is ‘acceptable’ or not. This latter purpose is served, in a large measure, by psychological and somatic signals generated by emotions associated with the beliefs.

Antonio R. Damasio has famously spoken of the role of *somatic markers* in the making of decisions and inferences in [30], mostly based on neuro-physiological and neuro-pathological observations.

The fact that inter-subjective and person-specific beliefs are often associated with emotions that act as *amplifying* mechanisms in the mind is of likely relevance in making some mental processes *sensitive* to initial conditions and, generally speaking, *unpredictable*, in spite of being deterministic. This is an issue we will take up later in this book, in chapter 9.

I close this section by emphasizing that beliefs in general, including the heuristics made use of in inferential processes, are mostly lodged in the unconscious mind, without the reasoning subject being aware of the action of these (incidentally, the commonly used connotation of the term inter-subjective implies a conscious sharing of ideas or thoughts, which means that our present use of the term is an extension of this common usage). In other words, inferences are predominantly carried out by ‘*hidden logic*’ where we have now stretched the meaning of the term ‘logic’ far beyond its commonly accepted connotation. More specifically, we will occasionally be using this term to mean any set of rules — whether inter-subjective or person-specific, and whether unconsciously or consciously invoked — that may be made use of in an inferential process. The set of ‘rules’, moreover, *may not be a consistent one* and may result in an inferential act involving a network of branches and cross-links, making it a tangled process with complex and contrary features.

The mechanism, outlined in the present chapter, relating to the possible role of beliefs and emotions in inference is in the nature of a speculation without direct empirical

support — though one that appears to me to be plausible enough so as to be accepted as a working hypothesis. It is not inconsistent with views of experts in the field of cognitive science as revealed, for instance, in comments in [112], chapter 1, where Stanovich, by way of quoting the social and political theorist Jon Elster, speaks of the distinction between ‘thin’ and ‘broad’ theories of rationality (the ‘rationality issue’ will be addressed in the next chapter). The so-called thin theory “leaves un-examined the beliefs and the desires that form the reasons for the action whose rationality we are assessing”, while, in reality, “we need a broader theory of rationality that goes beyond the exclusively formal considerations ... and that allows a scrutiny of the substantive nature of the desires and beliefs in action.” Here Elster (and Stanovich too) wants ‘desires’ to be included in an adequate appraisal of rationality, in addition to ‘beliefs’. I have not included desires in building up, in the above paragraphs, a hypothetical scenario as to how the inferential process can actually proceed in the human mind, where beliefs and emotions play a vital role (incidentally, desires are associated with emotions too). The inclusion of desires is essential when one goes on to examine the *values* governing the inferential process (I comment on this once again in the next chapter). This book pleads incompleteness in opting to steer clear of this all-important question of values in so far as these set the goals and purposes of inferential processes in everyday life and in scientific explorations. Instead, I look at the inferential process in the human mind as it proceeds with an aim to meet with *given* goals (*earning a Ph.D.*, or, *winning a war*).

Chapter 7

Reasoning and rationality: the intrinsic and the extrinsic

Epistemic irrationality

“Though this be madness, yet there is method in it.”

Starting from the mid-sixties of the last century, there has been a great profusion of investigations, journal articles, books, critical analyses, and polemics in the literature in cognitive psychology on the so-called *rationality* issue. What does it mean to say that, in making an inference, a judgment, or a decision, a person is being rational, or else, is deviating from standards of rationality?

Standards of rationality rest on evaluation in terms of sets of normative rules. An act of inference (or one of making a judgment or decision) made in conformity with a set of such rules then earns the accolade of being rational. The question, however, comes up as to whether the rules are of an *intrinsic* or *extrinsic* nature. The actual process of inference followed by an individual may involve the use of intrinsic rules specific to that individual or of inter-subjective rules that may differ from the rules invoked to evaluate his standard of rationality. In cases that the rules are of an inter-subjective nature,

they may or may not be *universal* ones like, for instance, the rules of *logic*. The rules that the individual makes use of can, moreover, vary from one inferential act to another similar one and, what is more, may not even be fully known to that individual, since they may operate *tacitly*. Indeed the 'rules' may be so diffuse as to be quite unlike what goes by the name of rules. This raises the question of the extent to which the process of reasoning and inference-making can be described and evaluated, by persons other than the individual drawing the inference, in terms of a set of explicit rules invoked for the purpose of evaluation.

All this is sometimes expressed by saying that *normative* accounts of reasoning and decision making may differ from *descriptive* ones. However, this raises a subtle residual problem: can the internal psychological process in a person making an inference or a decision be at all amenable to a descriptive account by either the individual concerned or by some other person attempting the description?

Before engaging with these questions of central relevance, I'll relate to you in a few sketchy paragraphs as to what type of norms the cognitive psychologists have had in mind in conducting their psychological investigations by means of tests conducted on individuals as also on groups of individuals in the context of various types of 'tasks' given to them. It turns out that the performance of the subjects of these investigations is, generally speaking, rather ordinary, indicative of a rather low level of 'rationality' of these individuals. This makes necessary a re-examination of the concept of rationality that these investigations seek to reveal.

The rationality issue: a brief overview

For the purpose of the present discussion, I will use the term inference-making to include such psychological processes as deductive and inductive reasoning, making decisions involving choices between available options, and making judgments as to the relative probabilities of beliefs in the face of available evidence. Evidently, the processes mentioned in this list are not sharply differentiated from one another, since there exists a good deal of diffuseness in the definition of each of these and, at the same time,

a good deal of overlap. For instance, it is only appropriate that the term ‘reasoning’ should refer to processes involving the use of *rules* to some considerable degree where the rules should be inter-subjective or, more preferably, universal ones like the rules of logic. By this token, making a deductive inference should qualify as an act of reasoning. An inductive inference, on the other hand, is essentially one of informed guessing, where rules are defined much less sharply and the term reasoning is of dubious validity, though an act of inductive inference does involve the *processing* of information. On the other hand, an inductive inference involves the formation of new beliefs or the revision of our judgment as to how probable our beliefs are when faced with one or more new items of evidence. In the same vein, decision making involves the making of a number of inferences where one has to weigh the consequences of choosing between a set of alternatives.

Let us consider deductive reasoning first since it is a form of inference where one expects that people will be most ‘logical’, and will conform to normative or evaluative rules to a large extent. Mathematical deduction is taken to be the ultimate form of deductive reasoning, but most psychological tests are performed with much more simple forms closer to our everyday exercises in deduction, one widely studied form being syllogistic reasoning where a few premises are supplied and the subject is asked to check whether a proposed conclusion follows from these. In most cases, the answer is uniquely obtained by the application of *modus ponens* (‘affirming the antecedent’) or *modus tollens* (‘denying the consequent’) or combinations of these, or else (in the case of more complex syllogistic tasks), by the use of *Euler diagrams* ([85], chapter 2), where these constitute the normative rules for this form of inference making. However, people do not always follow the normative rules (which, incidentally, are *universal* ones in that these do not involve inter-subjective rules of limited validity, or person-specific rules) and are commonly found to arrive at erroneous conclusions. In particular, errors are quite frequent in the case of syllogisms or logical tasks posed in *abstract* terms. Subjects perform better when the tasks are posed in terms of concrete themes and concepts that they can relate to, but then they appear not to follow consistently the valid rules of syllogistic reasoning, but to respond *contextually*, i.e., by relying heavily on their *interpretation* of

the tasks and by referring to their *prior knowledge, concepts, and beliefs*. For instance, they are found to exhibit *belief bias* ([81], [85], chapter 4, [125], chapter 11) i.e., they tend to avoid logically valid conclusions if these go against firmly entrenched beliefs of theirs and, conversely, favor invalid conclusions if these are supported by their prior beliefs.

Refer to our previous discussion of belief as an evaluative element in inference. Incidentally, a rule that says that people often make inferences in syllogistic reasoning by weighting these with their beliefs held in common with others (a fear of the number thirteen, for instance), itself constitutes an instance of an *inter-subjective* but *non-universal* rule since it is not sanctioned by universal rules of logic but is still found to be used commonly while assessing the performance of subjects in psychological tests.

In other words, as far as the psychological processing of information goes, people often convert deductive tasks to ones resembling inductive inference (recall our earlier discussion of ‘problem view’ and ‘process view’ in chapter 4) and arrive at conclusions that are fallible and are not strictly necessary consequences of the given premises.

The next type of inferential tasks we will consider is *decision making*. Recall that we have used the term inference-making in a broad setting, where it includes the making of decisions. Suppose you are in a gambling frame of mind and have two gambling options open, each with some specific probability of win and an associated pay-off in money terms, and you want to decide as to which of the two gambles to accept. In real life, we continually face such ‘gambles’ where we have to choose from a set of uncertain outcomes, with each outcome associated with consequences that may or may not be preferred by us when compared with the consequences of the other possible outcomes. In order to make a successful choice one needs to do some reasoning and inferring. More generally, most reasoning and inference-making tasks in real life involve a lot of implicit decision-making in that the sequence of information processing steps through which an inference proceeds require decisions to be made (mostly without conscious awareness) as to which of several alternate courses to follow in order to arrive at a

successful conclusion. Psychological tests on decision making, however, are performed on explicit decision making tasks such as a concretely specified gambling situation.

The normative theory one starts with in gambling tasks is the ‘expected value’ theory where one calculates the expected values associated with the various possible choices in accordance with their respective probabilities and value outcomes, and then makes the choice with the maximum expected value. However, for the sake of generality, monetary value is to be replaced with *subjective utility*, where the latter is based on a subjective scale of preference that may not be amenable to specification in quantitative terms ([112], chapter 2). Further, the probabilities associated with the various different choices cannot, generally speaking, be assigned specific values. Still, a normative theory can be postulated in terms of a set of *axioms of choice* ([112], chapter 2), where a subject can be said to make rational decisions if she consistently adheres to these axioms. Consistency with these axioms then constitutes the normative principle of decision making.

For instance, the transitivity axiom tells us that if choice A is preferred over choice B and, in turn, choice B is preferred over choice C, then choice A has to be preferred over choice C. While this appears to be unexceptionable enough as a requirement for rationality, actual psychological tests have shown that people’s choice is often stamped with *context* effects when it comes to the transitivity axiom or other equally logical axioms of choice. Thus, unknown to the cognitive psychologist, the question of choice between A and C may bring in subtle context effects in the form of associations from past experience that result in C being preferred over A. For instance, to take a hypothetical and unlikely example just for the sake of illustration, a person may have a preference of even numbers over odd numbers and, among odd numbers, a preference of those divisible by three over those that are not, as a result of which she prefers 12 over 9, and 9 over 5, but when it comes to the question of choosing between 12 and 5 she expresses a preference for the latter number for a completely different reason (perhaps in a recently held ‘lucky draw’ she was asked to choose between these two numbers, when the choice of the number 5 proved to be lucky for her). Once again, the normative principle may be overridden by or conjoined with beliefs held by people, and the choices that people actually make in psychological tests of decision making deviate notably from what is

assumed by cognitive psychologists to be the norm in decision tasks ([112], chapter 2, [125], chapter 11).

I will conclude by briefly outlining a similar deviation from the normative principles relating to *belief revision*. Do people correctly judge their beliefs to be true or false? Are the beliefs consistent with one another? Do people correctly update their beliefs when confronted with new evidence? In short, do the beliefs held by an individual correctly reflect the state of the world as perceived by her? All of these relate to what can be termed the *epistemic rationality* of a subject ([112], chapter 1). If beliefs are all we go by when making inferences (which, in the present context, includes the making of decisions), then a necessary condition for the inferences to be on the right track is that the beliefs should have a good measure of correspondence with the actual state of affairs ‘out there’, since the world at large is the ultimate source from which the beliefs are formed and, moreover, is the source of the ‘premises’ on which these beliefs are made to operate (many of the premises are themselves in the nature of beliefs). Indeed, beliefs are the end results of inductive inference in the context of novel or salient situations, which implies that an act of inductive inference can be interpreted, in a broad sense, as one of *belief revision by means of available evidence*.

The normative principle for such belief revision is made up of the *axioms of probability theory*, along with a number of logical deductions from it, notably the principle known as *Bayes’ theorem* — a mathematical formula relating the ‘prior probability’ of a belief or a hypothesis (say, H), to its ‘revised probability’ resulting from an ‘evidence’ (E) coming to light. The revised probability represents the degree of belief in H , given that E is true (i.e., is observed for a fact), and may be denoted by the symbol $P(H|E)$ while the prior probability $P(H)$ is the degree of belief in H without regard to E . The formula involves two other probabilities, namely, $P(E|H)$, i.e., the likelihood of E occurring as a consequence of H and, in addition, the probability $P(E|\sim H)$, i.e., the likelihood that E occurs as a consequence of any and every hypothesis other than H , the two taken together determining $P(E)$, the unconditional probability of E being true.

Epistemic rationality is said to obtain if a subject updates her degree of belief in a hy-

pothesis (I avoid the clumsy phrase ‘degree of belief in a belief’) in conformity with Bayes’ formula. The question that immediately comes up is, can one ever hope to assign reliable values to all these various probabilities? However, it has been argued that exact probability values are, in effect, not as important as *consistency* with the formula: what one needs to do in order to be counted as a rational agent is to be consistent with Bayes’ formula in a qualitative sense ([112], chapter 3; see also [105], chapter 5, where Rosenberg examines the question of rationality in respect of scientific hypotheses and theories), without regard to exact probability values. In order to test how rational a subject is, a cognitive psychologist provides a subject with values of the probabilities $P(E)$, $P(H)$, and $P(E|H)$ (or values from which these probabilities can be worked out), and asks him to estimate what the value of $P(H|E)$ would be. In large numbers of tests of various descriptions of this type, it has been found that *people deviate in substantial measure from the axioms of probability theory* (recall that Bayes’ theorem is a logical corollary of the basic probability axioms). For instance, subjects often neglect to consider the *base rates*, i.e., the unconditional probabilities of hypotheses and, at times, tend to place $P(H|E)$ on the same footing as $P(E|H)$. Generally speaking, a number of deviations from Bayes’ principle can be interpreted as implying that people assign a greater degree of significance to concrete evidence than to abstract statistical information which, in a sense, is a context effect over again. In other words, people pick out what is more salient to them with reference to their information base while ignoring information less salient in their perception. However, salience is perception-dependent and is determined substantially by the belief structure of an individual, which goes to show that the way we set about to revise our beliefs when confronted with evidence, is itself dependent on our existing belief structure.

One other major deviation from the basic logic underlying Bayes’ principle relates to ‘ignoring the alternatives’, i.e., not paying adequate attention to $P(E|\sim H)$, the likelihood of the evidence E being true under hypotheses other than the hypothesis H under consideration. In other words, we tend to revise our beliefs in a ‘dogmatic’ manner — while focusing on some particular belief H (which is more often than not made salient by a set of existing beliefs), we tend not to examine how likely it is that the evidence concerned

can be explained in terms of alternative hypotheses: we are *biased* by our existing beliefs.

All these deviations in probabilistic inference notwithstanding, humans are found to be pretty good detectors of frequencies of occurrence of various kinds of stimuli received as inputs from the environment, and of degrees of covariation between groups of stimuli (see [112], [103]), where much of this ability of frequency detection seems to have been acquired in the process of biological evolution, and is in the nature of an innate or unconscious one.

Significantly, the same Bayesian approach is commonly assumed to provide the normative basis of *how scientific theories are confirmed* by evidence coming up through experimentation and observation, because theories are based on hypotheses that are arrived at by inductive inference (one belonging to a special category, namely, abduction). The confirmation of scientific theories is a widely discussed and debated issue in the philosophy of science where there have been many attempts at working out a normative principle of confirmation. No such logically sound normative principle has been arrived at (refer to [105], chapter 5, mentioned above), paralleling the fact that no normative principle of belief revision can be formulated that describes with any degree of universality the actual inferential processes of individuals.

Rationality: misplaced notions

I will not dwell further upon how and to what extent actual inference-making by individuals deviates from normative principles made use of by cognitive psychologists in evaluating the level of rationality commonly achieved by people (see, for instance, [37], [36], in addition to references mentioned above) . Indeed, the entire complex of cognitive processes in individuals is heavily context-driven. For instance, the process of memory recall, one of great importance in inference-making, is biased by *relevance* which, in turn, is influenced by the immediate goal or purpose, and by *available* data, information, pieces of knowledge, associations, and beliefs (i.e., in other words, memory recall

is, to a large extent, a ‘top-down’ process), since an *unbiased search* of the entire ‘data base’ in memory constitutes a computational task of enormous proportions ([125], chapter 3) . Whichever way one looks at it, inference-making by individuals appears to be non-ideal and less-than-rational, having only a limited correspondence with objective normative principles supposed to be infallible guides to successful reasoning.

It is only rarely that people base their reasoning exclusively on deductive logic or axioms of probability, or axioms of choice, or strategies in game theory. Instead, they guess, mostly under situations of uncertainty, and make copious use of their beliefs and of an enormous store of heuristics — thumb rules, hunches, and clues — in arriving at inferences. And, remarkably, this entire complex process of inference-making is largely carried out by the *cognitive unconscious*.

Significantly, the process of inference making turns out to be, on the whole, epistemically authentic in spite of the deviation of the inferential processes from normative principles, as discussed above. This is to be explained by referring to the fact that the belief-driven inferential process is heavily influenced by past experience, in which past successes play a positive role as compared with past failures, and also to the fact that inferences are subjected to incessant tests and checks, this time against rules that are relatively more solidly grounded ones.

The large body of data accumulated over the years from rationality tests of varied types, has been interpreted differently by different groups of cognitive psychologists. One major interpretation involves the viewpoint that human inference-making is a fundamentally *flawed* process, biased by beliefs and heuristics, the latter, according to this point of view, being the vehicles that take human reasoning away from normative efficiency. In contrast there is an alternative point of view that there is no question of any fundamental flaw being there in human reasoning which is, on the contrary, *effective and adaptive*. Human rationality, instead of being flawed or biased, is *bounded* [110], [45] in the face of the fact that most real life inferential tasks are, computationally speaking, of unbounded complexity. Moreover, most inferential tasks are necessarily carried out

under uncertainty and on the basis of incomplete data where, in the specific context of any particular task, the probabilistic and choice-theoretic approaches do not make much sense, and guessing becomes a *necessity*. In other words, because of uncertainties inherent in most situations, and of limitations of available data, and also because of bounded rationality in the face of unbounded complexity of most inferential tasks, humans *have* to place a high value on being *effective* rather than normative, where the effectiveness is achieved mostly at the level of the unconscious, by means of acquired beliefs and heuristics that make possible *fast and frugal* reasoning [86]. Concurrently, the cognitive process draws upon *biologically evolved* capacities of the mind (refer to [103], chapter 3, mentioned earlier in chapter 5 of this book; we will again refer to ideas gained from *evolutionary psychology* in respect inferential processes in later sections).

Based on bounded rationality, the cognitive system makes effective use of its elementary inferential abilities (detection of patterns of covariance and similarities in environmental inputs) developed in a long evolutionary process — abilities that get diversified and sharpened in the developmental process of the individual, where communications with other individuals in a community sharing a common culture play a vital role. All this results in an *effective* inferential ability of an *adaptive* nature. The cognitive system of bounded abilities focuses on and makes use of the essential features of the inputs received from the environment, depending on its immediate goals, without entering into an unbounded search process when faced with a complex setting. Its belief system, along with emotional mechanisms of amplification helps in the ‘selection’ of an effective inferential path.

The less-than-ideal and bounded rationality is, in essence, *a strength* since it limits the search for relevance and meaning in the multitude of environmental inputs and makes the search effective. In order to be able to judge and compare, a system has to be able to make a ‘meaningful selection’.

It is now time to collect what we have come up with so far so as to put together a plausible framework of the way the inferential process can conceivably be carried out in the human mind.

In so far as an inference involves a processing of information in the mind, it is in the nature of a sequential application of rules — broadly speaking, rules of implication — where, however, the rules are not of uniform measure of generality. At the top of the ladder are the universal rules of logic and mathematics — ones that are independent of context (i.e., the state of the world within which the inferential process is carried out), including the specific context defined by the web of beliefs of the individual, with her associated emotional referents, and the more general context of inter-subjective beliefs. Then comes the less general class of rules that are of an inter-subjective nature — ones that are not specific to an individual, but are still based on shared beliefs of a class of inference-making subjects. And, finally, one has to reckon with person-specific rules that are relevant within the inferential process of the individual but have no broader significance. Among these, the inter-subjective and the person-specific rules are context-dependent — not only the belief-laden psychological context of the individual or of groups of individuals, but the more general contextual setting as well, the state of the world if you like, within which the inferential problem is defined, including its ‘inputs’ and the desired ‘output’.

The ‘rules’ (the universal ones and the inter-subjective and person-specific ones of a belief-laden, heuristic nature) operate upon beliefs that get involved in the various stages of the inferential process, and make use of the knowledge base, memories of past experience, and the tangled mass of other beliefs, along with their associated emotional links. In this, the inferential process proceeds in the form of a bundle of a tangled sequences of ‘information processing’ executed in parallel, eventually ending up in a ‘conclusion’ or an ‘output’ that is more often than not a more or less diffuse bundle of concepts and beliefs (either newly formed ones or ones produced by modification or revision of previously formed beliefs) itself.

And what is more, the rules are seldom of an *explicit* nature in that they cannot be formulated or described with any degree of completeness by the individual engaged in the inferential act, or by an external agency — by, say, the cognitive psychologist. The reason lies in two things — *first*, the person-specific or group-specific, non-universal nature of the rules and of the context defining the inferential process (‘personal knowl-

edge’); and, *secondly*, the predominantly tacit nature of the process, where most of the processing is carried out within the unconscious stratum of the cognitive mind.

When looked at within this complex setting, the mass of psychological data telling us of the less-than-rational status of human inferential act does not appear surprising because the norms or evaluative standards invoked for the purpose of judging rationality are at odds with the actual manner in which the mind engages in inference-making. It is within such a complex setting that the rationality issue is to be addressed for an adequate resolution. Any simplification by restriction to a less complex setting is likely to result in an incomplete and misleading idea of rationality.

This, in essence, is indicative of the naturalist approach to the rationality issue. However, the naturalist point of view cannot be limited to just the question of how the cognitive inferential processes actually occur in the human mind, because that, even when amenable to being addressed meaningfully, amounts to just a plain description. Of equally vital relevance are the issues of *sharing* of abilities between individuals and communities so as to *develop* and *improve upon* the currently existing inferential faculties, and of examining the goals and *values* underlying the inferential processes. All this refers to broader ethical and moral questions relating to human reasoning, inference, and decision making, and to the question of *responsible* decision in a social context. The naturalist approach has to address all these issues to be a really effective and useful one. These broader questions will not be addressed in this book because of essential limitations on the part of its author, due to which he feels inadequate for such an undertaking.

The dual-process theory

It is here that I want to draw your attention to a currently favored view of the nature of information processing in the human mind — the *dual-process* theory. This is a view subscribed to by a good number of exponents in cognitive science and is based on the assumption, consistent with a large body of literature, that human cognitive process involves two distinct streams — *process-one* and *process-two*. Of the two, the former is

a fast, non-deliberative, massively parallel and autonomous process, while the latter is a slow and deliberative one, predominantly based on sequential processing that makes copious use of working memory. Significantly, process-two is perceived as being capable of acting as the control of the cognitive process as a whole.

The question of control is a complex and non-trivial one, though. Commonly, the controlling action of cognitive processes is assumed to be exercised by the conscious mind. However, it is not unlikely that much of the controlling action is, in fact, executed without overt awareness [124].

The dual-process theory is fundamentally in the nature of a suggestive and fruitful hypothesis — somewhat in the nature of a metaphor — a prototype theory that provides a framework for interpreting and explaining quite a large body of observations and information. As in many other areas of cognitive science, it is a useful *heuristic* itself.

Process-two operates on the basis of rules of relatively more general scope such as the universal ones of mathematics and logic, and the inter-subjective ones that some relatively large group of individuals make use of while being *aware* of these as reasonably reliable beliefs, though not foolproof ones (a number of chess heuristics, for instance, or a few — too few? — of the heuristics that stockbrokers make use of). Process-one on the other hand, is based on ‘rules’ of a more primitive kind operating, in the main, in the unconscious recesses of the cognitive mind — the person-specific ones and the relatively more fallible ones of the inter-subjective type, *along with* heuristics of some proven worth. While process-two is executed with at least some measure of focal awareness, process-one is more tacit and is in the nature of a *default*. While the rules in process-two are, broadly speaking, those of *implication*, those in process-one are less distinctly so, being based, to a larger extent, on processes of *association*. Finally, process-one is of an *adaptive* nature, where it is required to be effective in achieving some useful purpose or goal in the specific environment in which it operates while process-two is a *deliberative* one where epistemic goals acquire relevance. In other words, process-one is geared to achieving what is referred to as *instrumental rationality*, in contrast to process-two

in which *epistemic rationality* ('what is the world like?') is of greater relevance ([112], chapter 1).

What is of great significance is the manner in which these two process types *interact* in the course of an inference. This, of course, is a matter of speculation though, once again, one that may prove to be useful and suggestive in some considerable measure. It seems likely that both the processes run simultaneously though process-two has periods of relative latency while, at the same time, there occurs a continual transfer of control to it when it evaluates, by means of comparatively secure rules of implication, some intermediate idea, concept, or belief produced in the course of activity of process-one, and passes on the 'result' of that evaluation to the latter which is then used as an ingredient in the further processing of information carried out autonomously of process-two. In this, process-two is likely to contribute its own quota of information processing (apart from evaluating and *justifying* what process-one has handed over to it), handing back a more 'finished output' to process-one.

Medawar speaks of "a rapid reciprocation between an imaginative and a critical process, between imaginative conjecture and critical evaluation" ([88], p 44) in the context of the so-called hypothetico-deductive method of science, championed by Popper. The hypothetico-deductive method presupposes a distinction between phases of 'discovery' and 'justification' in scientific exploration. The two, however, can be distinguished only notionally since they are but two aspects of the same inferential process. The dual-process theory speaks of a more intimate mix of the 'imaginative' and the 'critical' processes where it need not even be necessary that there be two separate processes involved. A more likely scenario is that the inferential process is a single complex one involving 'imaginative' leaps and rule-driven 'critical' sequences. The latter picture of the inferential process is not inconsistent with the one hinted at in this book.

The 'context of discovery' is seldom addressed and analyzed in logical and philosophical discourse, since it is assumed to belong to the realm of psychology. However, psychology is no exception to the 'logic' engendered by the operation of causal connections. The 'logic' of the cognitive unconscious may be something completely different

from the logic found in textbooks, but it does operate on the basis of ‘rules’ nevertheless — rules of *less than universal validity*. In particular, as repeatedly hinted at in this book, the rules may have deep moorings in beliefs and emotions of a personal nature. The apathy, found in conventional logical and philosophical discourse, towards an analysis of the context of discovery was questioned by Hanson, who looked for a ‘logic of discovery’, thereby contributing important insights into the process of abduction. Hanson, however, fell short of analyzing the deeply psychological roots of the logic of discovery [52], [53]. In particular, it is likely that the ‘evaluation’ of an inferential process is carried out, in a large measure, within process-one itself, through the involvement of neural aggregates operating as ‘pleasure centers’ in the brain. These neural aggregates are of early evolutionary origin but are co-opted, in a manner of speaking, in complex inferential and cognitive activities of the mind.

It is in this transfer of process control from process-one to process-two and back (the “restless to-and-fro motion of thought” spoken of by Medawar in [88], p 48) where one is likely to find the clue to the way the reasoning capacity of an individual evolves in the course of his or her developmental history, and to the remarkable difference in the reasoning capacities of individuals as also to the difference in *modes* of reasoning of individuals and across cultures.

To be more specific, there are, in all likelihood, *two* things involved here, one dependent on the other. First, the quantity and quality of the vast repertoire of beliefs and heuristics that one can draw from in the course of an inference. And, secondly, the manner and the frequency of transfer of control from process-one to process-two as mentioned above or, speaking in more general terms, the manner of *interaction* between the two processes. When one stops to think of it, there can be an infinite range of variation in the mode and efficacy of the inferential process, depending on the interplay of these two factors. But one cannot say more since one has to draw a line here before indulging in dangerously unfounded speculation.

Emotion-driven processes play a great role in leading to remarkable variations in the

reasoning capacities among individuals, and to the progressive development of the reasoning capacity of a single individual. Success in an inferential act is a potent emotional factor that inspires enhanced performance in subsequent acts of inference of related kinds. Success elicits recognition and acknowledgment from peers, and public approval as well — all of these intoxicating indeed. No less intoxicating is the *internally generated affect* relating to success, and the resulting sense of confidence, of which a likely consequence is a heightened ability to explore possible alternatives in an inferential process — picking up cues in profusion from environmental inputs and making remarkable use of heuristics made available by the cognitive unconscious under the amplifying action of favorable emotions. In other words, emotions have an amplifying psychological role that may produce dramatic effects in the development of cognitive skills.

Numerous exponents of the dual-process theory converge on the opinion that process-one is of more ancient evolutionary origin than process-two, and there also appears to be a widespread acceptance of the view that the cognitive system, especially the one involved in process-one inference, comes equipped with a considerably developed *tool-kit* for adaptive thinking that has emerged in the course of biological evolution. While the ‘tool-kit’ is of an innate nature, it gets *expressed* [29] in the course of individual developmental process — a process greatly influenced by the cultural environment of the individual. However, it is no less likely that the innate inferential capacity of evolutionary vintage is supplemented with capacities not originating in biological evolution but ones developing in a secondary process in the course of the inferential history of the individual, though this is once again a point where speculation is apt to become too nebulous.

In any case, it is not unlikely that the dual-process viewpoint is a pointer in the right direction, though it is still somewhat in the nature of a suggestive hypothesis. While numerous psychological experiments seem to support the viewpoint and appear to make the distinction between process-one and process-two reasonably sharp and well defined, evidence from neuroscience is less direct, since it seems unlikely that there exist

clearly demarcated neural aggregates devoted to the two types of processing. It may, for instance, turn out that the dual-process point of view will be supplanted by a more elaborate theory involving an integrated but complex processing system, with the inferential process circulating between various ‘depths’ of this system. But that, for now, is beside the point. What really matters is that the dual-process view is not in discernible conflict with known facts, is consistent with numerous psychological tests, and is a richly suggestive one, providing clues to a more enriched understanding of human cognition.

The distinction between process-one and process-two has an intriguing parallel with that between inductive and deductive inference, since it seems but natural to surmise that induction and deduction are based predominantly in process-one and process-two respectively. The observation that every inference has aspects of deduction and induction intertwined with each other is then consistent with the one that process-one and process-two are deeply enmeshed in everlasting interaction with each other as well. Put differently, the dual-process point of view provides an enriched understanding of inferential processes in general and of inductive inference in particular.

Finally, the dual-process hypothesis supplies an appropriate and enlightening perspective to the ‘rationality issue’. Commonly invoked norms of rationality are capable, in some measure, of evaluating the performance of process-two but not of process-one which occurs mostly outside the domain of focal awareness and makes use of hidden resources presumably replete with heuristics, beliefs, and emotional links of a deeply personal nature. There exists a proposal that instead of one single rationality, one has to think in terms of *two* types of rationality, as espoused by Evans and Over ([37], p 8): a ‘personal rationality’, and an ‘impersonal rationality’. In a manner analogous to the dual-process view, it tells us that the rationality issue is more complex than what it has been supposed to be.

I will wind up this section by pointing out that there is speculation in *two* directions here: speculation that cognition in general and inferential acts in particular involve two *streams* of processing of information at the psychological level — namely, process-one and process-two that have distinct and complementary characteristics; and specula-

tion that cognition and inference-making are based on person-specific resources and resources of an inter-subjective nature that have a kinship with *person-specific* ones. The former is a hypothesis shared by a good number of experts in cognitive psychology and is acknowledged to be a richly suggestive point of view that is likely to pave the way to a more complete and meaningful theory of cognition. The latter, on the other hand, is more of a speculation, but one that dovetails rather seamlessly with the dual-process view and gives it substance. But, at the same time, it increases the complexity of the problem of understanding human inference-making by one clear notch.

Complexity

Does it pay to make a problem more complex than it at first appears to be? Shouldn't one rather look for *simplifications* so as to reach at the heart of the problem? One does not really hope to have a compelling answer to this, as there is no compelling answer to the ultimately right *method* to be followed in science. We will have more to ponder over this question in chapters to follow. But one can safely say that it does not really pay to make a problem *too complex* by bringing in all the innumerable details that may conceivably be relevant to it, because too much of details stands in the way of building up a meaningful theory. However, the question of complexity is of relevance at a deeper level of theory building. When looked at from such a deeper level, the question of the relevance of person-specific psychological resources in inference-making does appear to deserve attention. Related to this is the question of a basic *non-determinability* in psychological processes that is likely to have a bearing on the issue at hand.

The unpredictable

Consider a dynamical system that is known to be of a *deterministic* nature, having some well defined causal rules of evolution of its own. Several questions at a deeper level now come up and demand attention. *First*, even if the system in question does have a well defined set of causal rules of evolution of its own, *can those rules be known to*

us? Can we fathom out what the rules are? Or, are there hidden substrata to the system that have an essential control over its behavior, ones that stand in the way of an external agency (the cognitive psychologist and the neuroscientist in the present context) fathoming out the rule(s) of causal evolution of mental states? The history of science provides us with a partial answer to this question that we will be presently having a look at. *Secondly*, how important is the role of the *context* in which the behavior of the system is observed — what is the role of the *environment* in which the system is made to evolve? As far as the course of an inferential act is concerned, this is linked with the first question posed above, in that the *internal* context based on hidden psychological resources of an individual does seem likely to have an essential role in the dynamics of the cognitive system and so, we will now have to concentrate on the possibility of the *external* context having a role in the cognitive process. And, finally, even disregarding the question of context and that of the rules determining the system behavior being hidden from us, does a *deterministic* system dynamics guarantee that the behavior of the system is *determinable*?

All these are questions of vital relevance for our purpose. Starting with the third of the above questions first, the answer is already known to be a big *no*. An overwhelmingly large set of dynamical systems obeying even simple dynamical laws of a deterministic nature are known to be non-determinable, where the lack of determinability arises from *sensitivity to initial conditions* ([11], chapter 9) engendered by an essential nonlinearity in the system behavior. There may, however, exist large domains of observation where this non-determinability does not appear to be of essential relevance. A very big chunk of physics, for instance, rests on the linear behavior of the harmonic oscillator and on the nonlinear but deterministic dynamics of two bodies interacting by means of the gravitational interaction, where results of stupendous relevance have been arrived at. At the same time, there exist well defined approximation schemes where weak nonlinearities of the oscillator and weak perturbations over the two-body gravitational interaction can be taken into account in arriving at remarkable results over an even broader range of observations, without any of the non-determinability coming to the fore. In other words, just the *possibility* of non-determinability does not necessarily

mean that the non-determinability is *relevant*. This, however, does not tell us the whole story. There remain entire areas of observation where the non-determinability does acquire relevance. While a linear oscillator behaves in a determinable manner under a periodic external perturbation, the behavior of a nonlinear oscillator is of a wholly different nature, being, under certain circumstances, *chaotic* in nature. A similar incursion of chaos occurs in the behavior of three gravitating bodies under quite realistic conditions. One finds here that the question of non-determinable (but deterministic) dynamics is linked up with that of the context-dependence of system behavior, where the effect of the environment comes to the fore.

Even for a system with a chaotic dynamics, many coarse-grained features of the system behavior may be amenable to description and explanation.

The essential role of the context

Speaking in general terms, there is, of course, nothing new in the observation that the environment influences system behavior. There does remain the possibility, though, of something deeply intriguing in a certain kind of context-dependence, of which the foremost example is that of a quantum mechanical system, where the idea of the system having some pre-determined value of an observable quantity regardless of measurements performed on it to reveal that value, is not a valid one. Here the very context of measurement determines which of a *possible set of values* comes up with what degree of likelihood. The case of the psychological state of an individual has an uncanny similarity with the quantum situation, though the context-dependence of psychological states we are now looking at has little to do with quantum indeterminacy. The psychological state of an individual is often indeterminate without reference to his or her external context (the internal context, on the other hand, is implicit in the very idea of a psychological state). Consider, for instance, a simple situation when an individual is to choose sequentially (choice between A and B, B and C, C and A) from among a set of alternatives. Does there exist a prior order of preference for the alternatives in the psychological state of the individual regardless of the context in which the choice

is exercised? Not necessarily, because, for one thing, he does not always conform to the axioms of choice, of which the transitivity axiom is one (refer back to [this section](#) for a brief outline). Consider, in a similar vein, a number of consumer items, with specified prices, available to a housewife in the market. But the mere availability with their price tags (along with a given set of internal beliefs, desires, and necessities of hers) may not suffice in uniquely determining her choice in a marketing expedition since the beliefs, desires, *et cetera*, find different expressions towards different sets of things and in different situations external to her.

The question of context in the making of inferences in general, and of making an inductive inference in particular, is a deep and broad one, needing reference to 'external' and 'internal' aspects of the context, as I indicated in chapter 2 (see section entitled [All observations and theories are contextual](#)) while referring to the contextuality of scientific concepts and theories. More generally, any inferential act is constrained and conditioned by external and internal contexts in a complex manner. The external context includes signals received from the environment that set the purpose of the inferential act in question, including the ones that trigger the process of inference. The internal context, on the other hand, is made up by the mental resources that are brought to bear in realizing the course of the inferential process.

What is important to realize is that these two aspects are not independent of each other and are deeply intertwined, and that neither can be made explicit in terms of a set of symbols having specific connotations. The external context, for instance, includes signals picked up by the cognitive mind — but, how does the cognitive mind decide which signals to pick out from among the infinite number of signals sent out from the environment at any given instant? Evidently, the decision depends on some kind of salience and relevance with reference to the purpose of the inferential process under consideration. But, salience and relevance are matters of *judgment* that, in turn, require a set of inferential processes where *internal* resources of the cognitive mind are activated. Moreover, a vast number of cues are picked up, at any given point of time, from the environment by way of *tacit* or implicit cognition. Some of these implicitly received inputs are made use of in the currently activated inferential process, while some others

are stored in some form or other for use on future occasions. Myriads of cues picked up from the environment (mostly unconsciously) are made use of by the cognitive mind so as to make up the store of *heuristics* that act as rules in the inferential process (see below), again, mostly at an unconscious level. The rules constitute a major part of the internal context of the inferential process.

The internal resources of the cognitive mind mentioned above, some of which are brought to bear in any given inferential task, thereby setting the internal context of the latter, make up a vast canvass. This includes the belief system of the individual concerned, her knowledge base and memory of past experiences and, over and above all this, her emotions and affects that play a crucial role in her inference making. Most of these internal resources are made use of within the unconscious strata of the cognitive mind, in the form of *rules* or, more generally, as causal links in the process of inference. Rules made use of in inference are referred to throughout this book in a variety of contexts. I summarize below a number of major considerations underlying the idea of inferential rules.

The rules of inference-making

The process of inference is not a random or sporadic one. It has a point of origin in one or more mental representations that constitute the premises of the inference, as also a second set of representations making up the conclusion. In between, the cognitive mind proceeds through a series of intermediate sets of representations, where the transition from one set of intermediate representations to the next can be thought of as being effected by means of *rules* that operate in a manner analogous to 'if — then' clauses used in everyday linguistic expressions and also in computational processes. In a more general description, the transition from one set of intermediate representations to the next may be said to occur due to a set of *causal* mechanisms.

Of course, one cannot rule out some involvement of sporadic processes in an inferential act, since the cognitive mind remains forever engaged in processes of an associative nature, with associations being established between elements of thought without any

CHAPTER 7. REASONING AND RATIONALITY: THE INTRINSIC AND THE EXTRINSIC

specific purpose. Indeed, this apparently aimless but incessant process of association may be assumed to result in a vast repertoire of heuristics, to be made use of in diverse inferential acts as the necessity arises.

A second qualification that seems necessary here is that the sequence of transitions in an inferential process need not form a well defined linear one where intermediate premises are passed through one after another. More plausibly, the intermediate representations form a clumsy tangle where the sequence branches out, backtracks, and is crossed with many sub-sequences that arise from 'trials and errors' on the part of the cognitive mind.

Finally, as I summarize below, the rules and causal mechanisms cover a wide spectrum as regards their structure, ranging from universal rules at one end to deeply person-specific ones at the other, where emotions and feelings make their presence felt. Possibly, some of the person-specific rules are so diffuse and implicit that it may be quite impossible to identify these as causal links. However, I adopt the position that the cognitive process is predominantly a causal one, though the network of causation, in all likelihood, is vast, diffuse, and enormously complex.

As mentioned several times in the foregoing chapters, the spectrum of rules acting as the causal links in an inferential process can, at a rough reckoning, be described as being made up of universal rules, inter-subjective but not universal ones, and finally, the person-specific rules. The universal rules include those of logic and mathematics and, in addition, rules that have gained universal acceptance by the force of experience and have been incorporated as knowledge (*if a liquid is heated, it will eventually turn into vapor*). While universal rules are inter-subjective, not all inter-subjective rules are universal, notable among which are the ones imbibed from some particular cultural environment (*old men are not to be relied upon; women are of a delicate mental make-up; the number thirteen brings in bad luck*). Finally, person-specific rules can be more or less transparent ones (*Mr. X has had a divorce recently; consequently, Mr. X avoids women*) or, at the other extreme, may have their roots buried in some remote past, perhaps within a mass of forgotten experiences (*memory of an unpleasant encounter*

with a certain person in long-forgotten past causing an avoidance of contact with persons with a similar facial appearance). The latter are of such a diffuse nature as to have, at times, only a vague or weak causal role. It is, however, a hypothesis worth entertaining that person-specific rules play a quite considerable role in inference making. On the face of it, person-specific rules can only lead to errors in an inferential act. However, at times these can be effective ones, being in the nature of heuristics formed out of cues picked up from the environment at varied points of time. For instance, a good physicist is likely to have a larger repertoire of heuristics of a mathematical and physical nature (perhaps being half-formed intuitive ideas generated from a minute study of some set of natural phenomena and also picked up from the scientific literature) than others around her.

The majority of rules are in the nature of heuristics and beliefs, mostly operative within the cognitive unconscious. These constitute a part of the *internal* context of an inferential act. Another large part of the internal context is constituted by emotions and feelings, many of which are associated with diverse beliefs of an individual. Together, these make for a vast complexity characterizing the inner mechanisms of the cognitive mind.

It was Ludwig Wittgenstein who famously pointed out that, speaking on logical grounds, rules can never be formulated in unambiguous terms — there always remains the problem of interpreting and understanding what a rule says, and any finite number of examples of application of a rule and any finite algorithm for expressing what the rule is meant to say, has to be ambiguous. Saul Kripke [75] has interpreted Wittgenstein's take on rules in the field of language. According to Wittgenstein and Kripke, there has to be a tacit inter-subjective agreement on the universal application of a rule, defined by means of *examples and algorithms*. While Wittgenstein's and Kripke's approach to the rule-following paradox can be described as being an essentially logical-analytic one, it points to the deeply cognitive roots of how people come to follow rules. Indeed, rules are always interpreted from examples or algorithms in an *inductive* manner which is why they are, from a strictly logical point of view, always ambiguous. We come

to make effective use of rules by interpreting them within some given context where the context exercises a constraint in effectively removing the ambiguity. In this sense, even the universal rules of mathematics are, strictly speaking, ambiguous, and make sense to a mathematician only within the context set by her prolonged and continuing exposure to rule-following practices of a very special kind.

Guessing at the inner mechanisms of a complex system

This finally brings us to the first of the questions listed above (refer back to section entitled *The unpredictable*). How completely can we know the inner psychological mechanisms that operate in the making of an inductive inference, or in the making of inferences in general, by an individual? Given the complexities of the external and internal contexts involved in inferential acts as outlined in the present chapter and in the last, the tenor of the answer is already evident — our knowledge of the inner mechanisms of an inferential act is indeed *necessarily* incomplete.

But, does this incompleteness imply that any attempt to guess at these hidden mechanisms is futile and doomed to failure? Can it be that the guess is destined to be fundamentally wrong? This is not the fate that generally awaits hypotheses and theories in science despite the fact that *all* of these are in the nature of guesses — some more solidly supported by evidence than others, and *all* are based on incomplete knowledge of the underlying mechanisms. Even when a hypothesis or a theory is incorrect, it is seldom *fundamentally* wrong since more often than not it does correctly capture some aspect or other of the workings of the system it is supposed to model.

The same goes for informed guesses relating to hidden mechanisms in inference-making. These do tell us something important about these hidden mechanisms in inference and in cognition, namely, that there exist processes and resources in a substratum of our psyche, as a result of which the commonly invoked descriptive and normative accounts of inference making miss out on aspects of essential relevance, and, at the same time, provide us with significant pointers as to what these processes and resources may be.

The hypothetical account of how the processes of inference-making (especially, the process of making an inductive inference), set out in the above paragraphs, proceed in the human mind, also point the way to a resolution of apparently conflicting points of view relating to rationality, i.e., to the integrated capacity for reasoning and for the making of inferences (including the making of judgments and decisions) in the human mind. It tells us that our reasoning is indeed biased, but also tells us that the bias does not make the reasoning fundamentally flawed since a bias may, *under circumstances*, be more productive than strict neutrality in the judging of evidence, and may have an adaptive value. It tells us, in a similar vein, that the use of heuristics in reasoning and inference is indeed not a logically sound practice but, at the same time, *may be an effective one*. It is the context that makes the biases and heuristics of the human mind lead to confusion and error and *also* the context that, otherwise, makes these appear as resources possessed of a considerable degree of effectiveness. The same hypothetical account also tells us that the task of formulating descriptive and causal principles of inference-making through successive stages of improvement is not a hopeless one, much as the task of describing and explaining the *world at large* is not a hopeless one, acknowledging, at the same time, that the successive stages are not likely to make a neat linear progression toward an *ultimate truth* since it may, in all probabilities, involve *sweeping* changes in perspective as, indeed, *all* scientific journey does. But, we will have to wait a bit before we take up the issue of the course of evolution of scientific theories in general.

Chapter 8

Abduction: theory in emergence

The enigma of hypothesis formation

**The abductive suggestion comes to us like a flash.
It is an act of insight, though of extremely fallible insight.**

Charles Sanders Peirce

Abduction is, in the main, the formation of *hypotheses*. Hypotheses, in turn, are needed for the purpose of providing *explanations*. Hypotheses, in other words, are germs of *theories*.

Abduction has been described as ‘reasoning in reverse’ (Johan van Benthems foreword in [40]). Knowing that your wife has gone out, you can explain the crying of your newborn baby, who is dissatisfied with the attentions of her nanny, and also the absence of the familiar sound of the radio playing, both of which you detect before entering home on returning from your office. On the other hand, supposing that your wife has gone out before having had an opportunity of informing you beforehand, you reason ‘in reverse’ to explain the sound of crying of your baby and the absence of the sound of the radio playing, and come up with the hypothesis that your wife has gone out, so as to explain these two observations of yours. The hypothesis subsequently proves to be correct, but it was arrived at by guessing, since both the observations of yours could have had some

other causal origin (perhaps your wife has fallen asleep, or perhaps she has upset the baby for some reason *and* the battery of the radio has run out).

This is an instance of abduction, though not of theory in emergence. Here there is no revision of belief, only the inessential addition of new belief, and no new rule of operation on the beliefs. This is in the nature of an elementary hypothesis-making; the newly formed belief will shortly attain the status of knowledge, and will be stored as such, perhaps in the short-term memory. Hypotheses in science are vastly more complex and more intriguing.

The making of hypotheses lies at the heart of abduction, and hypothesis is *informed guesswork*, where the vast background of knowledge, beliefs, and emotions is made use of.

“A dominant further factor, in solid science as in daily life, is hypothesis. In a word, hypothesis is guesswork; but it can be enlightened guesswork” ([102], p 65).

The method of science has been described as the ‘hypothetico-deductive’ one ([88], p 45), where there is a constant interplay between processes of forming hypotheses, deducing the logical consequences of the hypotheses, testing these consequences against facts of observation, and revising the hypotheses to ones with a greater explanatory power (the acts of testing and revising are based on what are commonly referred to as the ‘logic of confirmation’ and ‘logic of refutation’). This seems to be a fairly good description of the scientific method, though the question of ‘explanation’ and that of the ‘scientific method’ are contentious ones. While the actual practice in scientific enquiry does not always fit with this scheme of things, it provides us with a convenient *notional framework* for discussing various aspects of the scientific process. For instance, as indicated by Medawar, a number of salient features of the scientific process are explained in terms of the hypothetico-deductive scheme.

The issue of *explanation* in science has been a widely and critically examined one, with a large body of literature devoted to it. A big part of that literature is analytical

and logical in nature, dating from Hempel and Oppenheim's influential paper [58] that introduced the so-called Deductive-Nomological ('D-N') model of explanation. Hempel and others subsequently revised the model in several stages, but the D-N model gradually lost appeal, as it was realized that explanation has more aspects to it than can be subsumed under a monolithic logical account. For instance, the relation of explanation and causation has for ever been a contentious issue. Pragmatist, naturalist, and cognitive-psychological approaches have been tried so as to have a coherent understanding of the idea of explanation, where elements of the D-N model also find their place, and a substantive understanding now appears to be within sight (for an introduction, see [64], [129]). In this book, however, I will go only by the *common-sense* meaning of the term 'explanation'.

The topic of 'scientific method' has also been a hotly debated one. We will have a brief look at it in the next and the final chapters of this book. In particular, the hypothetico-deductive method is crucially dependent on the idea of *confirmation* or corroboration of scientific theories where confirmation, in turn, relates to the judgment as to how and to what extent a theory reveals *truths* about nature. Theories of confirmation, however, quite frequently run into troubled waters precisely because scientific theories are arrived at principally by the route of inductive inference (see [28] for the idea underlying the hypothetico-deductive approach; see also [109], chapter 5, [105], chapter 4, for background). The hypothetico-deductive method, while of notional value, does not describe science as it is actually done.

The generation of hypotheses: an adventure into the mysterious world of psychology

The hypothetico-deductive scheme as conceived by Popper and others, while acknowledging the psychological origin of hypotheses, actually gives it a wide birth owing to the fact that the psychological process of generation of hypotheses cannot be probed by philosophy and by logic.

“[Philosophical views of hypothesis generation] All aim at demonstrating that the ac-

tivity of generating hypotheses is paradoxical, either illusory or obscure, implicit, and not analyzable.”, [82], p 1.

The naturalist approach to philosophy, on the other hand, takes a hard look at what scientific practice itself has to say of this psychological process of hypothesis generation, where cognitive psychology makes a close scrutiny of this process as it is found to occur in individuals, and itself forms *hypotheses* as to the inner mechanisms underlying this process. In this, cognitive psychology gets a good measure of help from *artificial intelligence* where various aspects of hypothesis formation are realized by means of computer programming. At the end of the day, though, it remains to be seen whether the process of hypothesis generation, while being, in the main, a causally determined one, is also fully *determinable*.

Hypothesis formation is, of course, not something that is specific to science. The formation of hypotheses in some form or other, is a necessary adjunct to *every* act of inductive inference since, in the latter, the ‘conclusion’ or the ‘output’ is never a *necessary* consequence of the ‘inputs’ one starts from. An inductive inference is rarely, if ever, a one-shot process where one straightaway forms a conclusion by some divine inspiration, but is one that has to go through several stages of reverse reasoning where hypotheses are formed and discarded, though one is not usually aware of such trials and errors (at times, the cognitive mind takes advantage of past trials and errors available to it in the form of recalled experiences and of relevant heuristics). With abduction being thus associated, at one end, with mundane acts of inductive inference of everyday occurrence, it assumes the form of the remarkable and awe-inspiring phenomenon of *creativity* at the other. In between, it provides the scientist with her staple food of the intellectual variety, namely, *hypotheses* of various degrees of plausibility from which emerge the *models* and *theories* in the sciences.

Abduction and novelty

In other words, abduction covers an extensive spectrum of psychological acts, where there is a corresponding spectrum of the degree of *novelty* associated with an act of

abduction. When a child gleefully toddles to a nearby cupboard to retrieve a toy that her mother had previously hidden as a challenge, and then jumps to the mother's lap for appreciation, she has definitely done something novel, as the mother is always the first to appreciate. And this first success in meeting the playful challenge of the mother transforms the world of the child (and that of the mother too!) as only a novelty can do. Then again, as a Ph.D. student forms a hypothesis regarding the mathematical calculation of the bond energy of a complex molecule which later is proven to be a useful one in the context of her dissertation, she has achieved something novel, that her supervisor happily appreciates. These are, however, not commonly appreciated as being *remarkable* except, perhaps, by a few intimate observers because the transformations resulting from these novelties do not cover a 'conceptual space' of a wide span. What brings out a really ecstatic and overflowing appreciation from a great number of people is that supreme creative act of a van Gogh or an Einstein. Leaving aside, in the context of this book, the creativity of van Gogh, that of Einstein is reflected in the observation that, as the theory of relativity emerged in the form of hypotheses in Einstein's mind, the latter impacted on the entire edifice of theoretical physics to send a great number of previously entrenched concepts topsy-turvy and paved the way to a host of new ideas and concepts.

The idea of the *conceptual space* is used here in the same spirit as in [11]. Roughly speaking, it stands for the entire set of concepts in the mind of an individual. However, concepts are *interconnected* thought entities, and the conceptual space is endowed with a rich *structure*. The whole of the conceptual space is, strictly speaking, of little relevance since what makes more sense in any given context is some sub-space (in a loose sense) made up of concepts meaningfully related to the context. In an act of producing a hypothesis of some novelty, the relevant part of the conceptual space is *restructured* and *enlarged*, with new combinations of earlier concepts being formed and, in addition, strange and new concepts getting included in an expanded conceptual space. It is the multitude of concepts in a conceptual space that get combined into beliefs, some of which get transformed into knowledge.

While the idea of conceptual spaces is used in this book in a rather loose and intu-

itively defined manner, these are referred to in more concrete terms in various contexts in artificial intelligence, cognitive science, and in semantics and linguistics (see, for instance, [42]).

Briefly stated, when our beliefs and working rules of inference, including those of rigorous logic, fail to make appropriate predictions or to explain a body of evidence, a hypothesis is called for wherein a new combination of beliefs, or an altogether new set of beliefs is formed, and a new rule of operation in our conceptual space emerges (more commonly, however, a new evidence is explained with the help of the existing belief structure and existing rules, by way of adding beliefs by simple combination or modification). Now, what we refer to as ‘new’ or ‘novel’ is, in a sense, only seemingly so since the germ of the new must have been there within the womb of the old. Thus, new life is created from existing life by the dissociations and combinations of molecules, in accordance with chemical rules operating in the world of those molecules. The creation of new life is a novelty principally because we are not fully aware of the actual mode of operation of these rules and of the complex dissociations and associations that take place (I do not mean to make a prediction as to whether these will eventually be known at some point of time in the future), and the process continues to elude conceptualization on the basis of the existing conceptual structure even after great efforts. Scientists know a great deal about the process, but they cannot *re-create* it.

Paralleling the process of creative hypothesis-making by an individual, there takes place the process of conceptual change and emergence of an entire new theory in the collective mind of a scientific community. In the former process the hypothesis emerges by way of receiving a certificate, so to say, from the ‘internal censorship’ within the mind of that very individual, perhaps operating unconsciously. In the case of a scientific community, a new theory emerges and is accepted as one in a complex process. There are deep parallels in the two processes, since neither the emergence of a hypothesis in the mind of an individual, nor the emergence and acceptance of a novel theory in the collective mind of a scientific community is a simple or a straightforward process. Neither is driven by logic alone. Both depend on the belief system,

mostly operating, in one case, within the unconscious mind of an individual and, in the other within the *substratum* of the collective consciousness of the community. The substratum in the latter case is constituted, precisely by *individuals and groups* within the community under consideration. The beliefs and desires — even conscious ones — of individuals and various groups within the community circulate within the latter as latent resources in its collective consciousness. Once a hypothesis emerges that impacts upon an entire conceptual edifice, it is acknowledged as a novelty by larger groups of people, and even by the person or the community that produced the hypothesis in the first place.

In other words, the novelty is as much in the process itself as in the surprise and lack of comprehension of the beholder. Einstein's thought process did not cause as much of a precipitous change in his own conceptual world as it shook the conceptual world of the physicists of his time and it is this, more than anything else that made his hypotheses so novel. And this, in turn, relates to the fact that the mechanisms underlying the formation and emergence of the hypotheses were, by their very nature, not explicit to the physics community at large and perhaps not explicit, in a large measure, to Einstein himself. Most of these mechanisms operated in a *substratum* of his mind, where a large number of heuristics and beliefs relating to ideas in physics and mathematics went through a prolonged process of association and dissociation, of feverish activity and relative latency, the exact nature of which may not be fully known to cognitive psychology in days to come. And it is fundamentally the same hidden nature of mechanisms operating in a substratum of the human mind that lies at the heart of novelty in the child's discovery of the hidden toy and in the Ph.D. dissertation where a new method of calculating the bonding energy of a complex molecule is presented.

The child's feat, on the other hand, involves primarily a restructuring of her own world of possibilities of action and thought and is a novelty in so far as her mother cannot imagine how this could come about — the child herself has little awareness of novelty. Her unconscious feeling of novelty finds expression in a sense of joy in her new-found ability.

There is, of course, a commonly (and, informally) perceived *scale* of novelty. The child's feat is appreciated as novel by only a few other than her mother because such feat is observed routinely in almost all children. The Ph. D. dissertation is also a rather routine one unless it gives the community of chemists something that has a sweeping efficacy. But a restructuring of almost the entire world of theoretical physics is rare indeed, and is perceived as a mind-boggling novelty.

In other words, novelty has two aspects to it: a transformation in the structure of a conceptual space whose mechanism is not explicit but is hidden within some unobserved substratum, and the extent and impact of that transformation. Added to this is the potential (and subsequently realized) ability to cause a further expansion of the conceptual space. A successful hypothesis is indeed possessed of novelty in this latter sense — it brings in new concepts, a new way of making sense of those concepts, new explanations and predictions, and an expanded scope of making sense of the world around us and, over and above all this, is obscure as to its mechanism of emergence.

Abduction: the naturalist point of view

This hidden mechanism of the emergence of a hypothesis is precisely what the naturalist approach to the philosophy of science is interested in. While concrete investigations in cognitive psychology and in artificial intelligence struggle to understand aspects of this mechanism bit by painstaking bit, naturalist philosophy picks out cues from the results of these investigations to make *generalizations* so as to add meaning to our *yearning* (a *fundamental* yearning of mankind, if you like) to understand this world of ours, as also the world within ourselves, wherein epistemological questions in philosophy appear in a new light.

It is this hidden mechanism that the dual-process hypothesis of human cognition and the hypothesis relating to role of heuristics, beliefs, and emotions — mostly entrenched in the cognitive unconscious — in inductive inference in general, aim at making a little more intelligible, and a little less obscure.

Abduction: the crucial link in science

Abduction constitutes the crucial link between cognitive psychological processes taking place in the hidden recesses of the mind of the individual and that stupendous product of the collective intellect of *humankind*, the response of Mankind to the mysteries of Nature, that 'legend' we call Science. It is crucial not only because of its importance and relevance to the scientific process, but more because it is enigmatic in nature, and mankind does not know how exactly the conjectures of science well up from within its own bosom. It is a helpless admirer of its own product of awesome fecundity, unable to reproduce it at will or even to comprehend it by means of logic or philosophy.

Abduction, in so far as it constitutes the most crucial link in the scientific process, concatenates, on one side, conjectures and hypothesis coming up, mostly, from within the substratum of awareness of the individual, and collective deliberative powers of entire scientific communities on the other. In this, abduction may be said to constitute the link between the *context of discovery* and the *context of justification*.

The two terms have been greatly discussed ones in the philosophy of science, where 'context of discovery' refers to the psychological process by which novel ideas are arrived at while 'context of justification' refers to the process where an idea is tested against facts of observation and also against other ideas. Reichenbach, the noted philosopher of science, is commonly credited with introducing the two terms into the discourse relating to scientific method. Historical and philosophical perspectives on the distinction between the context of discovery and the context of justification can be found in [108].

The process of abduction: exploration of a conceptual space

The context of discovery lies within the mind of the individual, made up of the enormously complex mosaic of knowledge base, combined with the vast repertoire of heuris-

tics, beliefs, emotions and motivations of that individual, and is truly and specifically, *personal*. The individual, on being motivated, by means of stimuli coming from the world at large, in relation to some problem that does not find a solution within the conceptual framework of existing theories ('how to explain the regularities of the hydrogen spectrum'), looks at it from the vantage of this context, and lets it digest within the chemistry going on within the substratum of her own mind where there takes place an extraordinary churning of all the ingredients of inferential activity. It is in this sense that abduction is a *top-down* process, like much of cognition at large, and differs from a 'bottom-up' one that is principally driven by continual perceptual inputs. To be sure, the mind does draw in conceptual inputs (and perceptual ones too, especially in the case of experimental science), but these do not *drive* the process of abduction on their own strength, becoming effective mostly by first being integrated within the contextual mosaic of the mind and probably providing the latter with a fresh supply of heuristics. The inferential activity proceeds within process-one, but is continually referred to process-two as well where, in both the processes, there take place an effective 'censorship', preventing the inferential exploration from venturing too capriciously.

More specifically, inferences in general and abductions in particular, have two aspects to them, as we have seen on earlier occasions in this essay. One is the aspect of an apparently spontaneous process of association between ideas circulating in the mind — once again, mostly in the unconscious hinterland in it — one that never ceases and that goes on producing new ideas, new hunches, new beliefs, even without any overt guiding purpose. And the other is that of the 'censorship', and the harnessing of the spontaneous process with the help of *rules*. The vast set of rules is mostly of a person-specific nature, but also has impersonal elements (inter-subjective and logical ones) finely woven into it where most of the rules are, strictly speaking, *beliefs* of a specific type, some of which are *relevant* in the context of the problem at hand and impart a *direction* to the inferential process. How the cognitive mind makes an effective whole of these two contrary aspects of the inferential process, is the great mystery that cognitive science is to explore and unravel in days to come, though a fundamental element of non-determinability or unpredictability is likely to come up in the process.

The fundamental importance of the role played by the association between thought elements in the process of making of inferences constitutes a basic idea in cognitive psychology that I have based on in this book and forms part of what is broadly referred to as ‘associationism’, which has had a long history with a great deal of critical appraisal both for and against. An account of associationism is to be found in [84].

Amplifying mechanisms, instabilities, and parallel explorations

We have had a brief look at the inferential process in chapters 6 and 7. The process of hypothesis generation is but an inferential process of a special kind. Inferential processes, in turn, constitute a component of cognitive processes in general that take place within the overall dynamics of the human mind. And this dynamics is of a *computationally complex* nature. It is fundamentally nonlinear, made up of feedback loops, and has an immense number of relevant dimensions to it. Even as one tries to restrict the dimension by way of posing the problem of inference and of abduction in more or less specific terms, the reduced problem, in order to be meaningful, is likely to inherit the computational complexity of the human mind at large. For instance, there is a great likelihood that the dynamics, even restricted to the problem of hypothesis-making is possessed of ‘sensitivity to initial conditions’, the essential pre-requisite of non-determinability, being symptomatic of an essential element of *instability* in the dynamics. Thus, a psychological process driven predominantly by rules proceeds along some specific trajectory in the conceptual space and is, on the whole, a *stable* one in that it does not deviate appreciably from that trajectory. On the other hand, a process of inductive inference continually branches out from the stable rule-driven trajectory and follows a course of a *globally exploratory* nature because of *instabilities* being operative from time to time at certain junctures on a rule-driven trajectory. Once a trajectory branches out along new courses, there subsequently appear fresh rule-driven trajectories, again of a stable nature, till the occurrence of further branching(s) owing to instabilities. An instability is essentially an *amplifying* mechanism that aggravates an initial deviation from a stable trajectory (‘sensitivity to initial conditions’) whereby that deviation, instead of being smothered and

corrected, eventually results in a globally distinct trajectory in the conceptual space.

In other words, a rule-driven process of inference is of a step-by-step local (more specifically, locally *stable*) nature, where, at any point of time, the next step in the process is determined by the application of some specific rule(s) on the results of only a few earlier steps (typically, on the *current* psychological state of relevance) and, in this, the rules may fail to lead to any result based on the resources available in the current state. It is here that resources *other than* the currently operative rules become relevant, these being, precisely, beliefs and heuristics that are, in the context of the current psychological state, relatively *remote* ones. The reason why the psychological resources of a remote relevance get involved in the inferential process is that *associative* processes never cease in the human mind, and spontaneously and incessantly bring together resources from remote recesses within the conceptual space. As an impasse or a psychological gap is arrived at, at some point on a rule-driven trajectory within the conceptual space, these associations create new possibilities in the form of *local perturbations* over the trajectory, one or more of which may now be *amplified* so as to result in a fresh trajectory or a set of trajectories along which the inferential process proceeds once again on the basis of rules where, perhaps, one or more new rules have now become operative.

Inferential processes are, thus, of an immense vitality, where a major contribution to this vitality comes from the operation of relevant *emotions and affects*, because these emotions and affects are among the basic mechanisms underlying the processes of *amplification* within the human mind (interestingly, emotions and affects also play a stabilizing role when needed). When, at critical junctures in an inferential process, the need of fresh conceptual explorations arises, a likely thing to occur is that a set of relatively remote beliefs and heuristics become involved by means of spontaneous processes of association, some of which *acquire relevance* and introduce a local *instability* through a set of associated emotions and affects, thereby initiating fresh branches of rule-driven trajectories. The affects can be, broadly speaking, of two types — positive and negative ones, of which the former give rise to the instabilities while the latter have the effect of smothering of perturbations whereby the inferential process continues to retain its *local* character. In other words, certain beliefs, with their associated emotions, assume

the role of *evaluators*, either 'approving of' or 'frowning upon' perturbations around a trajectory.

An inferential act is an exploration. It consists of an exploration of the conceptual space made up of ideas and concepts by proceeding along a trajectory in that space, where the trajectory is generated by the operation of rules of various types. However, the conceptual space is one of an immense size, with a large number of dimensions to it where these 'dimensions' are linked with the diversity of the external reality that is captured by the cognitive mind. In any act of inference- or hypothesis-making, the conceptual space has to be necessarily truncated to some manageable size where only a limited set of relevant ideas and concepts make up the space to be explored, but even so a rule-driven trajectory of a simple type may not be effective in the exploration of that truncated space. There are two aspects here that need to be examined. One is the effective 'size' of the truncated space that is to be explored, while the other is the efficacy of the process by which that space is actually explored.

Evidently, the effective size of the conceptual space that needs to be explored in guessing at the reason why the baby is throwing a tantrum is quite negligible compared to the size and the 'dimension' required for guessing at an explanation of the black body spectrum. In cases such as the latter, the need may, moreover, arise whereby the 'size' and 'dimension' of the space one starts with has to be substantially increased midway during the process of abduction since the earlier size proves to be inadequate in producing a hypothesis of requisite explanatory power. It is here that the idea of *proliferation* of the inferential trajectory in the conceptual space comes up.

Trajectories proliferate in the conceptual space by mean of local instabilities whereby a single rule-driven trajectory branches out into more than one trajectories of a similar nature but where new sets of rules and a larger repertoire of beliefs and heuristics are brought into the process of exploration. The reasoning process, in other words, fans out along parallel trajectories (where the cognitive mind engages in *counterfactual* thinking, or thinking by imagination), some of which may later be found to move further and further away from any satisfactory solution to the problem at hand while others

may run along more favorable courses and may, in turn, give rise to further episodes of branching, and so on. In trying to form an effective hypothesis the mind starts with some general idea as to what is needed in order to arrive at a satisfactory explanation of the phenomena in question or, in other words, what, generally speaking, would constitute a solution to the problem at hand. For instance, it may be necessary to find a formula that fits a graph, or a number of graphs, obtained from experimental observations, where a formula obtained from an existing theory is found to deviate considerably from the graph. If, at an intermediate stage of exploration, a tentative hypothesis results in some formula that deviates from the graph to an even larger extent than the initial one, then the line of reasoning leading to that intermediate hypothesis has to be abandoned or kept in abeyance while, on the other hand a better fit with the graph would be indicative of an acceptable line of reasoning.

A mundane act of inference-making or abduction involves a conceptual space of relatively small size, and the trajectory resulting from a line of reasoning needs to fan out to only a small extent before a reasonably acceptable hypothesis is arrived at. To be sure, even such a mundane instance of abduction involves *some* episodes of branching and some vaulting over logical gaps, and *this* constitutes novelty of an embryonic type. But further along the scale are those acts of hypothesis building in which local instabilities and branchings occur in abundance so that all the parallel branches along which the reasoning process proceeds make a highly effective search of the conceptual space and also transforms the conceptual space itself by drawing in concepts and ideas that were earlier left outside the confines of the truncated space the process started with. The newly inducted ideas, beliefs and concepts then enter into the never-ceasing process of association and eventually cause further proliferation of the process of exploration, ending up, in rare instances, in a veritable implosion and the emergence of a radically novel hypothesis.

The question of scientific creativity has been analyzed and examined from various different angles in [11], [106], [83].

Abduction: the setting up of new correlations

The association between fleeting ideas and the resulting involvement of some of these in the inferential process may be looked upon as events whereby *correlations* are established between seemingly remote ideas, beliefs and heuristics with reference to the goal or purpose (*finding a hypothesis that will lead to a formula fitting a graph*) towards which the process has been set in motion. It is conceivable that this process of setting up of correlations within a conceptual space may feed upon itself, resulting in a grand correlation centered around some nuclear conglomerate of relevant concepts, whereby a remarkable conceptual structure of stupendous novelty emerges for the whole world to see and admire. In this, the reasoning process gets transformed from the local exploration of a limited conceptual space to a highly correlated global exploration of a greatly enlarged one where, in all probability, the rules of exploration get transformed in stages to more and more efficacious ones (perhaps approaching rules of more and more general applicability). Such a process, resembling a *phase transition* in a macroscopic aggregate of molecular units (such as the transition from a disordered gaseous state to a much more ordered liquid state) would then count as one of abduction located at the other extreme of the spectrum of novelty as compared to a mundane act of abduction of daily occurrence. The spectrum is indicative of the fact that, while acts of inference and abduction performed during the daily routine of an individual involve logical leaps in common with creative abductions of great significance in science, the latter is a process of rare occurrence indeed, involving the setting up of correlations between remote ideas and concepts wherein these ideas cohere to give rise to a hypothesis of remarkable novelty. In this, a great scientific hypothesis resembles great poetry or great music where remote, varied, and disparate feelings and emotions are brought together into a creation of supreme beauty that, at the same time, is elusive too.

What is spontaneous and elusive in a scientific hypothesis produced in a creative act of abduction in the mind of an individual is subsequently captured, at least in some tangible measure, in the scientific practice of an entire community of scientists where the latent possibilities in a hypothesis are now realized by deliberative analysis, by the detailed working out of ideas, and by means of experimentation and observation.

The drama now moves over, in a manner of speaking, to the arena of *the context of justification*.

The reasoning individual: an atom in a cosmos, and a product of evolution

The individual is not an isolated atom removed from the rest of the universe. Howsoever withdrawn she may appear to be while a hypothesis brews in her mind and she confines herself within the cocoon of her own mental world, she is but a microcosm of society at large, encapsulating within herself an enormously large set of greatly varied inputs from the latter in the form of beliefs, ideas, and ideologies, most of which are tacitly acquired without awareness on her part. Within this matrix of ideas and beliefs imbibed from the world around her, she is acutely focused — again, perhaps, without being fully aware of it — on some of the unresolved conflicts and problems in that world, among which one or a few belonging to a common family may set her mind, principally the unconscious hinterland of it, to frenetic activity. Why and how some particular unresolved problem activates the mind of an individual is again a matter of psychological chemistry, where her prior developmental history that determines her existing body of beliefs, motivations, and psychological propensities acquires relevance. Given the right fit between these propensities and some particular problem that ‘appeals’ to her, her cognitive mind, aided by the amplifying role of her emotions and affects, takes over. The rest will, perhaps, be history.

Along with ideas, beliefs and motivations imbibed from the world around her, the individual inherits a large number of inclinations and capabilities of a psychological nature from the unfolding history of *biological evolution*, of which she is a product. These appear as *innate* psychological capabilities in her. The processes of spontaneous association of ideas of an embryonic nature, and a number of embryonic ‘rules’ of inference are, in all likelihood, part of such inheritance. Further rules of inference and a vast repertoire of heuristics are subsequently added in the course of the developmental history of the individual under consideration.

The parable of the two children

Here is a hypothetical, though, perhaps, not wholly unfamiliar, example of how a 'sensitivity to initial conditions', and a resulting local instability may play a determining role in the developmental history alluded to above. Imagine two children of nearly the same age growing within the common environment of a family and receiving similar training and instructions. One day the two are asked to solve a number of arithmetic exercises when one of them receives a bit of extra praise and appreciation from a number of elders in the family on her performance while the other, by default, felt slighted in some measure. This seemingly insignificant event may trigger a course of development of far-reaching consequences, when the first child engages feverishly in mathematically oriented activities and self-training, eventually developing a great mathematical ability, while the second child shies away from mathematics and becomes indifferent to exercising his mental faculties. This can be taken as an apt metaphor for what actually happens in the mental world of an individual because here the role of emotions in amplifying the response to an event of psychological relevance is in evidence.

Analogy: the great organizing principle

I will wind up this chapter by way of dwelling upon a number of factors that often play a great role in inferential processes of individuals, especially in processes of abduction resulting in the formation of scientific hypotheses. Of these, the foremost is the idea of *similarity*. Sensing the similarity between ideas and concepts belonging to disparate parts of the conceptual space is one of the great faculties of the mind that has, in all likelihood, evolutionary roots, and that is capable of playing the role of a uniquely efficacious organizing factor in a conceptual space. The psychological ability of sensing of similarities between perceptual inputs (*a certain habitat somehow appearing to be similar to one of an earlier experience, where a predator was discovered hiding*) is, of course, possessed of a distinct adaptive value for individuals and may therefore have been fixed as a heritable trait in the course of evolution of species. As this faculty gets enhanced in the course of the developmental history of an individual by means of explorations of an inferential and abductional nature, where the ability of sensing of similarities be-

tween perceptual inputs gets transformed into the ability of sensing similarities between *conceptual* constructs and then blossoms into a unique psychological ability, it creates *correlations* between concepts residing in far-flung recesses of the mind, and thereby enlarges the conceptual space explored in the course of the formation of a hypothesis, making a coherent whole of a great number of apparently disparate ideas. In the process, new ideas get woven into the conceptual edifice of earlier vintage, which is now pruned of a number of obsolete and burdensome concepts belonging to it. The new-found conceptual structure is now the embryo of a theory, to be taken up and developed by an entire community of scientists.

One of Douglas Hofstadters great aphorisms: "...analogy is the fuel and fire of thinking" [60].

Further principles: simplicity and elegance

The other organizing factors of relevance are *simplicity* and *elegance* of hypotheses. The faculty of developing hypotheses having a quality of simplicity is again, arguably, of evolutionary origin where it facilitates the economical production of efficacious hypotheses and saves wasteful exploration of the conceptual space. The idea of elegance is, in all likelihood, a derived one but is capable of effecting a correlation within a conceptual space and also an enlargement of that space, whereby new ideas are inducted in that make a coherent whole with ones belonging to the limited space of earlier vintage. However, the beneficial role of the *belief* that ideas of simplicity and elegance have an inherent ability to generate useful hypotheses leads to the further belief that *the mechanisms underlying the external reality* — to which all hypotheses are ultimately directed — *are also of an essentially simple nature*. Such beliefs are inductive inferences at a higher level, made use of in the processes of abduction in diverse specific areas, and *may not correspond* to the external reality at large. In other words, the useful organizing principle of hypothesis-building, to the effect that hypotheses should be simple and elegant in keeping with the nature of the reality that they pertain to is, in all likelihood, of limited relevance. A belief in this principle is formed and developed in the

course of inferential processes conducted within external contexts of relatively simple structure, and may not be of a deeper relevance, being inadequate in contexts having a computationally complex and prohibitive structure.

Abduction: the inspiration of science

In summary, acts of abduction cover a wide range inference-making, with a correspondingly wide spectrum of the degree of novelty associated with those. Of special relevance to the scientific process are the generation of hypotheses in the minds of individuals that can be either of limited significance in the field of science (*a Ph.D. thesis whose findings do not have a wide impact*) or, on the other hand, can be of such great relevance as to initiate a new theory with a novel conceptual structure. Such processes of creative abduction involve, in all likelihood, a highly efficacious mode of exploration of a conceptual space where a cascade of local instabilities, resulting from amplifying mechanisms associated with emotions and affects, lead to a proliferation of inferential branches running in parallel and where there occurs a transformation of the conceptual space, with new sets of concepts and new rules of operation in an enlarged conceptual space getting involved in the process. At the same time, there appears a correlation among seemingly remote concepts and beliefs, as a result of which the concepts acquire a new coherence. An organizing principle of great relevance in bringing about such correlation and coherence is that of *similarity*.

Simply stated, *abduction constitutes the element of inspiration* in the progress of science.

Chapter 9

Summing up: science as an interpretation of the world

The complexities of truth

Many reflective people read with a sense of bemusement the obituaries so frequently (and so gleefully) written for scientific realism.

Philip Kitcher

Scientific theories pertaining to the workings of Nature have two interfaces: one with the minds of individual scientists and of communities of scientists, and the other with — Nature itself. At the first interface, theories are produced and given a finished form, since individuals produce hypotheses that are transformed into finished theories in the hands of communities of scientists. And, at the second interface, a theory looks out to the world and gives us a perspective view of what lies out there — how wheels move within wheels.

Scientific theories: the basic issues of relevance

We have focused on a number of aspects of the processes of inference and abduction taking place within the minds of individuals, these being predominantly ones of unconscious cognition where beliefs and emotions have a large role to play and where the mind essentially engages in an act of informed guessing, while at the same time relying on concomitant processes of testing, checking, and justifying. Further processes of development, testing, and justification of theories are carried out by communities of scientists where numerous individuals communicate, compete, and co-operate — finally producing finished theories that find general acceptance by the scientific community. We will have more to say on this aspect of scientific theories in this chapter of my book, before I summarize all our considerations in the next, final chapter.

The question that comes up at this point is: what kind of picture a finished theory draws of nature? How reliable, complete, and accurate can one expect it to be? In what sense can a theory be considered to be an improvement over previously developed ones, and whether a theory can be expected to be so complete and perfect that it can be described as forming a part of the ultimate frontier of human knowledge? All these pertain to the second interface that scientific theories come to possess, and will also occupy our attention in the present chapter.

As for the effect of the scientific community, and of society at large, on the scientific process, there is once again an entire spectrum of opinions. At one extreme, there is the point of view that social exchange among scientists and socially monitored scrutiny of the intellectual output of the individual scientist purges the latter of the stamp of individual beliefs, attitudes, and idiosyncrasies, and produces theories of a doubly objective character — objective in the sense of independence from subjective notions of the individual, and again, in the sense of being a correct and faithful representation of the mechanisms of nature. And, at the other extreme of the spectrum one finds the view that it is the society at large that essentially constructs our theories of nature by mechanisms insidious and pervasive.

The business of science is to *explain* — explain whatever is at odds with existing and established hypotheses and theories, and then to develop applications of the concepts and theories resulting from the process of abduction, giving rise to explanation over some specific area of inquiry. In this, science has to seek for *truth*, for unless it catches a true picture of some slice of nature, it cannot successfully produce explanations and make use of the explanations in developing applications of these. This entire matter of inferring ‘truth’ from success in achieving explanations, and of the sense in which a successful explanation adds to our knowledge of the world, is fraught with subtle and conflicting issues of philosophy, both ontological and epistemological. The ontological questions all relate to the basic one of what the inferences of science are true *of* — are these inferences and hypotheses true of a mind-independent world, or do these inferences provide us with just a set of effective *means* of operation in a phenomenal world that is the only world we can meaningfully talk about, while all talk of the *actual* world is, simply, so much of idle speculation?

The epistemological questions, on the other hand, revolve around the veracity of the statements that science offers us. Since, in arriving at inferences, the mind operates on *representations* of the world out there, everything depends on how these representations are arrived at. There is an undeniable gap between the representations and the real world that cannot be bridged in any foundationalist way with no vestige of doubt. This brings in a fundamental skepticism towards the theories of science that is very hard to shake loose. One such issue of skepticism questions the attitude of science towards *unobservable* entities and theoretical constructs, asserting that ‘truths’ about those should not, in all honesty, be accorded the status of genuine truths since these are, at best, nothing more than instruments of manipulation in the world of concepts. Moreover, the truths of science are *defeasible*: they are subject to revision from time to time in the light of new evidence or a new conceptual framework. What sort of truth, then, does this speak of?

Indeed, the very *rationality* or, one can say, the very *rationale* of science is not immune to challenge. Does science have a well defined aim? Has there been, in the history of science, a discernible progress towards the fulfillment of that aim? In what sense

do the successive theories and the successive waves of conceptual change in the world of science bear testimony to a cumulative progress made possible by the *methods* of scientific exploration? Or, to be yet more insistent, can one at all speak of a specific set of methodological *hallmarks* to be found in scientific endeavors, in contrast to those in other areas of human interest?

Realism or anti-realism?

Volumes have been written on questions alluded to above, and many more will be written in days to come, since none of these questions admit of a clear-cut answer acceptable to all. None of the attempts at answering these can be likened to an exercise in arithmetic, since the way one approaches questions of this sort depends on one's *point of view*. One can be a realist or an anti-realist, and even that does not describe one's point of view completely, since there are so many shades of realism and so many shades of anti-realism. To be sure, the exchanges between the realists and the anti-realists are not sterile or wasteful, since these help to throw a number of basic questions in sharp light, and that light may help one to *choose* a point of view, since a point of view is not really a bad thing when one sets out on an intellectual exploration.

The term 'anti-realism' has gained currency since the work of *Michael Dummett* who was principally concerned with the question of *truth* (we will address, in our own way, the complex issues relating to the idea of truth of scientific theories in section entitled *The two facets of truth* below). Can it be said of truth that it either does or does not apply to a concept? For instance, is it or is it not true that the world exists independently of our mind? While the realist would say that we certainly *can*, the anti-realist would like to understand first as to whether it carries any meaning to ask such a question: is there a way to *know* whether the world exists independently of our mind? This controversy about truth reaches out to far-flung areas of discourse. For instance, the philosophy of *intuitionism* in mathematics, to which Dummett subscribed, questions that of the Platonist-realist position regarding the truth-value of mathematical entities.

However, a point of view is not so bad only so long as one *acknowledges* it as such and does not set out to *prove* his own point of view as being the right one. Because, *that* starts the acrimony. For instance, you cannot ever prove that there exists a real world out there, to anybody who is bent upon looking the other way, simply because it is not a matter of pure logic but one of *experience*, and experience is no logic. Different persons sum up their experiences differently. In the course of our experience, each of us makes very many inductive inferences of various kinds and magnitudes. The existence of a world out there regardless of my thoughts is a matter of one such inference, a huge generalization (though seemingly a ‘natural’ one), one of an enveloping nature, but a non-deductive inference nevertheless. You can say that this inference has helped you in innumerable ways in making further inferences and in moving ahead in life, but that still does not make it an irrefutable one. ‘Very well, then’, you say — ‘give me a refutation’. But I *can’t* do that, you see, because this issue has *not* been settled to everybody’s satisfaction — not ever, nor can it ever be.

What is more, the response to the basic ontological question need not be a simple yes-no type answer. For instance, one may admit the existence of *entities* observable or unobservable, but may raise doubts about the *theories* pertaining to those entities, since inference to the *existence* of entities is more a matter relating to scientific *practice* than to theory. Theory building, in this view, is to be distinguished from the actual manipulation of natural entities where the latter, rather than theoretically postulated *mechanisms*, are of more direct relevance. As we see below, it is indeed of some use to distinguish between the question of the existence of natural entities — some of which we can observe more directly than others — and that of the concepts and theories abstracted from actual existence. This is an issue of some complexity and is related to the one I address below: to what extent can truth be attained regarding the mind-independent world, assuming that one exists. At the same time, as I briefly mention below, the two issues of existence of entities and of the theories of the way these are correlated with one another, can be separated from each other only conditionally.

Two books widely referred to in this context are [18] and [50].

Linked with this basic question of the existence of a mind-independent real world, there arise a lot of other questions, some of which I have hinted at earlier in this chapter. But here a further set of considerations assume relevance. Once again, the answers to these other questions depend on one's point of view, but now there arises the tendency of answering these within a *dichotomous* framework — with reference to one's position on the big basic question that lets one think in terms of two distinct *camps* having contrary orientations ' the camps of realism and anti-realism. For instance, granting that there exists a real world, is it possible to progress in a cumulative manner towards a true and complete description of that reality and of the mechanisms inherent to that reality? This is a question of understandable relevance, to which I doubt if a final or ultimate answer can be found, though I do have a point of view, as you may also have (incidentally, my take on the big question is, yes, there does exist a world out there independent of my mind or of anybody's mind; indeed, our minds are part of that world). However, this question of the possibility of progressive attainment of truth about the world needs a distinct domain of discourse as compared to the one relating to the existence of a mind-independent world. In other words, even when one subscribes to a positive answer to the basic ontological question, one may legitimately raise doubts about the nature and quality of *truth* that can be attained in describing the mechanisms of the world and about the way our knowledge about the world advances through successive waves of conceptual transformations.

This relates to the question, raised above, as to whether one can admit the *existence* of unobservable entities of science while, at the same time, denying the legitimacy of *theories* about those entities. Indeed, the question of existence of these entities is deeply linked with that of the theories about these. For instance, the existence of electrons cannot be de-linked from its properties that tell us how it interacts with other entities in this world, where the latter, in turn, is linked with theories describing the electron and all these other entities. Of course, it is possible to imagine that *something* answering to the name of an 'electron' exists regardless of our current theory of particles and their interactions. For instance, the concept of the electron is likely to survive in the event of a possible revision of the *standard model* of elementary particles (the concept was there

even before the emergence of the standard model). In a bigger context, this is analogous to a partial independence of the question of the existence of a mind-independent world and that of our theories of the world — mind-independent or not.

Theories: from the immediate to the remote

Theories about the world do not come in hermetically sealed packages delivered by divine intervention, but are built up in a long and arduous process of cognitive quest, where the quest starts with our mundane *day-to-day* interactions with the parts of the world immediately accessible to us. The inferences arrived at in the course of that interaction are primarily in the nature of inductive ones — tentative, fallible, open to revision and improvement from time to time, but *effective* nevertheless. The process of our cognitive quest in the accessible parts of the world is punctuated with events of a substantial conceptual transformation in one specific domain or another. For instance, we learned that water does not exist in the liquid phase under all conditions of temperature and pressure, and that realization must have resulted in a big transformation in our conceptual framework while, now that this fact has been incorporated within our conceptual framework, it does not make us wonder even for a moment. On the other hand, the question of the properties of water near the *critical point*, which is a much more specialized one, has been a matter of very definite concern to the experimentalists and theoreticians, and have been the subject of a deep and radical theoretical-conceptual transformation regarding phase transformations. In other words, mundane experiences and generalizations make way to progressively deeper theories in any domain of experience about the world. Put differently, the journey from the immediate to the remote in the world of concepts and theories is, in a sense, a continuous one and, at each stage of the journey, the cognitive system builds upon what is available and apparent to what is relatively more remote and inaccessible.

Kitcher explains in [72], chapter 2, how Galileo convinced his critics of the existence and motion of remote heavenly bodies observed by means of the telescope. His logic was essentially based on the idea that there was nothing in principle that distin-

guished between the sighting of a remote earth-bound body and a heavenly body that could be observed with the naked eye, and similarly, between the latter and a remote heavenly body that could not be so observed.

Put differently, theories are transformed in successive waves of conceptual transformation, starting from our mundane day-to-day concepts, and get built up in a hierarchical process where conceptual layers are built one upon another, and where ‘truths’ pertaining to one conceptual layer get revised in the succeeding layer. But, a truth that gets revised is no truth at all, and the currently accepted ‘truth’, that resulted from the revision of the ‘truth’ belonging to a former conceptual layer, will itself get revised in days to come and hence, likewise, is no truth at all. This appears to destroy the very idea of truth.

The two facets of truth

The term truth carries two implications — an implication of inter-subjectivity, and one of pointing to a perceived (or perceivable) fact of the world. Within a given context, truth, broadly speaking, refers to a perceived fact, where the perception (with or without the aid of relevant instruments) does not depend on the perceiving subject (within reasonable limits of variability of individual perceiving power). As the context changes, the fact may appear in a new light and may be perceived differently. Suppose that a certain biochemical reaction is found to occur at a constant rate within a given range of physical conditions. Suppose further that the course of the reaction depends on the temperature of the medium in which it is made to take place, but that fact was not known or suspected at an earlier state of knowledge. The constancy of the rate of reaction was then a truth in that state of knowledge. However, imagine that it has subsequently been discovered that at a much higher temperature the course of the reaction becomes oscillatory. The truth has now become *more complex* in nature. And, in the future it may assume a yet more complex form when the currently undetected dependence of the course of the reaction on some other physical parameter is discovered and recorded. The successive forms, differing from one another, in which scientific truth appears also depend on the

currently existing theoretical framework and on the history of its successive transformations. For instance, scientists, on observing the transition from a constant to an oscillatory behavior in a class of mechanical systems, may be in an exploratory state of mind as to the possible nature of the course of a chemical reaction, which may influence the way they set up their experiments.

The above example is, perhaps, a simplistic one, but it tells us that ‘truth’ is a complex concept indeed. It is, at times, asserted loosely that truth is a relative concept, but that does not mean that truth can vary from person to person (since if such a variation is indeed found to occur, then it has to be symptomatic of some peculiarity of circumstances that needs separate investigation). Truth, to be more specific, is relative in the sense of depending on the context in which facts of observation are recorded, but not in the sense of being dependent on the perceiving subject.

The context in which the truth of this or that scientific theory is judged keeps on changing, where the context itself is a concept of some complexity. For instance, it includes both internal and external factors with reference to which a fact of the world is perceived — internal ones relating to the current state of concepts and theories within a given domain, and external ones to conditions and circumstances under which the fact are obtained. Let us consider the external context first, focusing on the domain of observation, or the *horizon* within which a theory is constructed and applied. Mendel's theory of inheritance was confined within the domain defined by the phenotypic characteristics of species, where the molecular unit of inheritance could not be accessed or investigated within the then current scope of observations, and it was this that set the horizon, or the context, in which the theory was constructed. The conclusions of the theory, which were true within the context, were all *reformulated* in terms of genes and DNA in the altered context of cell biology, chromosomal studies, and molecular biology. None of Mendel's conclusions were *literal* truths when looked at from a more recent perspective and, by the strict, non-contextual definition of truth, these were no truths at all. Nevertheless, these were truths indeed since, by the then current standards of knowledge and reasoning, these did not require any *personal* belief to be interpreted and judged — anyone willing to perform a series of breeding experiments could see the veracity

of these truths. To be sure, the *origin* of the theory was within the personal confines of Mendel's mind, albeit with inputs from contemporary thought and from painstaking observations over a long period, and the inferential processes responsible for his theory were most likely of a strange and deeply personal nature, as we have tried to guess at in chapters 6, 7, and 8. But once formulated, his formulation had the stamp of a genuine scientific theory precisely because it was now *open to all* to be verified by experiments of diverse design and to be checked for consistency with the then established body of concepts. This is all that a theory needs in order to be accepted as true, and this does not require truth to be non-contextual or immutable. What is specifically interesting in this context is that, the transition from Mendel's theory of inheritance to the later day genetic theory involved a sea change in the relevant conceptual framework, where the true statements of Mendel's theory lost their *literal* significance in the new framework, but could nevertheless be *interpreted* as ones being of genuine relevance in the latter.

I will come back to this issue of the context-dependence of scientific theories and truths, since this is of great relevance regarding the question of the nature of scientific *progress*. For now, however, I switch to the question of the extent and manner in which scientific theories *refer* to aspects of the real world. Once again, informed discourse on this question needs subtle and complex considerations but I will be brief and direct here.

A scientific theory is a representation of a part of the world, arrived at by cognitive processes. The true conclusions of the theory are not literally true since these are inferences within a representation that *refers* to some part of reality but is not a *facsimile* of that part itself. It is somewhat like a portrait of a lady, where the portrait is a representation of the lady but not the lady herself. Features of the portrait refer to or, are *indicative* of, features of the lady but those features are the result of an *interpretation* on the part of the artist, and the relation between the features of the lady and those of the portrait is not a simple, geometrical one.

There is possibility of some confusion here that needs be clarified: the lady herself is perceived by the artist as a representation in his mind, while the picture is an expression of that representation in a different medium. Thus, the representation

of the lady in the portrait tells us (but again, in a skewed manner) how the lady is represented in the mind, only partly conscious, of the artist. Apart from the artist himself, there are three parties to the story: the lady herself, the representation of the lady in the artists mind, and the the expression of that representation in the portrait drawn by the artist.

In representing reality, we make categories, comparisons and distinctions. The question thus arises as to whether analogous schemes of categories, comparisons and distinctions apply to reality itself. In other words, does the representation have a relation of exact correspondence with reality, on the strength of which one can define truth? Can one, in the language of mathematics and cartography, make a one-to-one mapping from the lady to the portrait, setting up this correspondence? This is a question difficult to answer since 'reality' is accessible to us only through the representations, and we are never sure as to what the 'features' of reality are. But the relation between a theory, which is a consequence of a representation, and the part of reality under consideration, has to be, in *some* sense, one of correspondence since, otherwise, the theory could not have been successful in explaining relevant issues pertaining to reality or in developing relevant applications. This is once again, the argument of *success* and *effectiveness* implying *truth* in a theory, where 'truth' can be interpreted as a correspondence, in some sense that cannot be made fully explicit, to features of reality.

Features of the portrait of the lady in question do indeed correspond to some of those possessed by the lady herself, but the correspondence is strange and elusive. A portrait done by a great artist does strongly resemble the person painted in it, but the strength of the resemblance does not always lie so much in the physical features of the person portrayed as in what is commonly termed her 'personality'. It is the personality that complements the physical features of the person which the artist captures with an uncanny vision — a vision in which the inner self of the artist reaches out to that of the person portrayed, with a great sensitivity. In this sense the portrait is an *interpretation* of a person by the artist, where its features *refer* to those of the person in strange and elusive ways.

Representations in science are of a similar nature. The scientist, in making a hypothesis, gets into a strange mode of thought where her inner world reaches out to a part of the outer world in an intimate communion and she finds a vision not shared by people around her. This, of course, refers to the complex cognitive processes occurring within her mind, principally in hidden depths of it, the final outcome of which is a hypothesis subsequently accepted and worked upon by her peers. The hypothesis is eventually made into a theory that is now a representation of a part of reality — distinct from the internal representation within the mind of the scientist, but a representation all the same. The representation is essentially a *model* (in this context, see [43], chapter 1; according to Giere, the model corresponds to reality in the sense of an analogy) of that part of reality that the scientist and her peers are interested in (this sets the external context of the representation) that refers to and resembles the reality much as the portrait does. But, at the same time, the model is elusive and skewed because it has its moorings in the *internal* context defined by un-articulated beliefs and modes of thought of individuals who, in the final analysis, are microcosms of a community of scientists and, in a broader sense, of an entire culture. This is what constitutes the *dual* aspect of a hypothesis and, eventually, of a theory and, indeed, of any interpretation of reality. Every child, in the course of her development, makes millions of such interpretations, makes millions of predictions, based on anticipations, without being aware of those. Most of these predictions fail in small ways, but many others succeed, and this course of failures and successes continues throughout her subsequent developmental history.

A theory enjoys only a skewed fit with nature principally because it is a *contextual* description of the latter, and necessarily ignores *anomalies* lurking at remote corners bearing testimony to hitherto unrecognized aspects of natural reality. Equally responsible for this are the modes of thought of an entire scientific community, and its beliefs tied to the existing body of knowledge and the existing theoretical framework.

In other words, the duality of ‘truth’ marks all our representations of the world, ranging from the mundane representations in day-to-day activity, right up to the most rigorously thought out scientific theories, where imagination mixes with judgment, intuition mixes

with deliberation, and induction mixes with deduction. The difference between a 'trivial' representation and a complex and elaborate one lies in the exact degree and manner in which these contrary aspects intertwine and interact, giving a unique flavor to each individual representation conceived by the human mind.

The dual nature of representations gives rise to enigmas and questions. How can a representation, which involves aspects of imagination and intuition, be true of the reality it is meant to represent? Is the truth of a representation a relative one? Is the representation, infected as it is by beliefs and modes of thought of individuals and communities, irredeemably committed to giving us a *false* conception of reality, regardless of *where* that reality is situated?

While theories can indeed be false, those that survive and propagate have truth imprinted in them — a truth that is inter-subjective and not relative, but one that is contextual — a truth that, at the same time, is conditioned by the modes of thought of individuals and communities.

If a number of schoolchildren be asked to describe their school building, the descriptions given by most of them will be true representations of the building but will still be different from one another in the way they make their representations. The difference will lie not only in the literary styles and the modes of expression, but also in little ways in the *content* of each representation. This, of course, means that truth, even though inter-subjective in nature, is a complex thing. When asked what they find in front of them, each will say that she finds the school building and, at *this* level, there is absolute agreement in the 'truth' they speak of. But when it comes to a question of describing the building, their descriptions will differ in the *aspects* of the building they focus on. And, this difference reflects the difference in modes of thought of the children. Depending on the internal context defined by belief systems and modes of thought of individuals and communities, different aspects of reality are represented differently in scientific theories as in the descriptions offered by the schoolchildren.

To *summarize*, truth is objective (in the dual sense of being inter-subjective and of

referring to a common reality) but is contextual, and is *conditioned* by an underlying subjective world. In other words, truth, a feature of supreme relevance in representations, is a concept of exquisite complexity, incorporating within itself deep tensions between contrary facets. The huge turmoil, conflict, and confusion suffered by human beings in the pursuit of truths of this world are therefore in the nature of a pre-ordained and unavoidable destiny, for ever throwing men into depths of despair. Truth is not something *given* to us that one can sit back and discover passively, but is something one has to work for, struggle for, and stake one's soul for.

Anti-realists lose no time in focusing on the subjective aspect of the representations embodying scientific theories, and assert that the theories are socially *constructed* ones. And, realists, in a spirit of dichotomous thinking, emphasize on the aspect of truth, often underplaying the aspect of *interpretation*. But the *truth* about truth is more complex, as I have tried to tell you, and one needs a picture of science where this contrary and complex nature of truth is included. Of course, the tension between the contrary aspects of truth shows up in the works of philosophers of science, many of whom express their concern on this dual nature of truth. Realists do accept that truth is not a simple and given feature of a theory, and acknowledge that there is a sense in which one can say that truth is 'constructed'. And, anti-realists do acknowledge that there is something like truth in a theory in the sense of the *effectiveness* of the latter. In other words, realism and anti-realism are two camps *set up by us*, the great classifiers that we are, by force of a dichotomous mode of thought to which we appear to be committed in life, and the borderline between the two camps is not an indissoluble one.

The descriptions offered by different schoolchildren of their school building are *skewed* with respect to one another, and also skewed with respect to the building itself as it stands as a piece of reality, and this very skewness is a symptom of the fact that each of the offered descriptions is an interpretation generated in the mind of an individual student, which is thus somewhat of a *misfit* when considered in relation to reality — a representation that *refers* to and *corresponds* to a part of reality, but is still not a facsimile of it. The correspondence is strange and exotic in the sense that it is conditioned by the mode of thought peculiar to some particular student or other. The case of a

scientific theory is of a similar nature.

The two-faceted, tension-ridden and complex nature of truth raises deep questions for both realists and anti-realists, and people perceiving themselves to belong to one camp or the other face genuine problems in coming to grips with what exactly it is that the truth of scientific theories refers to. A case in point is Poppers attempt at explicating the idea of scientific truth in terms of the concept of 'verisimilitude' that he introduced ([90], see also [25]). The idea of verisimilitude, however, is more of a formal nature, seeking to capture the complex process of achieving truth in a succession of stages, than one in line with the cognitive-naturalistic approach where, moreover, one recognizes that the search for truth is contextual and *episodic*.

Hypothesis and theory: the individual and the scientific community

Here we have to contrast the aspects of interpretation and construction inherent in a *hypothesis* to those in a *theory*: the former is conditioned by the mode of thought of an *individual* (or a closely-knit group of individuals) and the latter by that of a *community* the scientific community that examines and accepts a hypothesis, and develops it into a theory. We have had a glimpse into the process of inference and abduction in the mind of an individual where beliefs and emotions of that individual have a role to play. However, that role is, in a sense, an indirect and supportive one since the beliefs and the associated emotions make possible an effective conceptual search and conceptual re-organization without themselves being used as ingredients in the formation of the concepts and conclusions in the inferential process. It is true that the process makes use of heuristics that are in the nature of beliefs themselves, but those heuristics are ones that are constantly tested against evidence and against previously tested concepts, and are revised sequentially so that the final conclusions that emerge, and the conceptual reorganization that accompanies the inferential process, have a measure of confirmation, justification, and truth built into these. The stamp of the individual that

remains in these final inferences and concepts consists of an overall conditioning of the latter by the mode of thought specific to that individual. For instance, a scientist may be concerned more with analytical-deductive thinking than with intuitive-explorative one, which directs her inferential process along a course different from the one that she would have adopted in case she had a different mode of thought.

The processes of conceptual organization realized by the inferential activity of a community of scientists have an analogous character. In trying to understand this, one has to appreciate that the scientific community, in turn, is a part of a society at large just as the individual scientist is immersed within the psychological and intellectual ambience set by the group or community within which she works. The analogy extends further in so far as the scientific community can be thought of as being a *cognitive system* itself. The inferential activity of the individual scientist has a deliberative component and a component of an intuitive nature, where the latter occurs mostly within the unconscious sphere of cognitive activity. This unconscious cognitive activity is, moreover, based on massively parallel processes in a network of independent neuronal aggregates, where the communications between the aggregates lend a great variety to the processes actually taking place in the network ([21], [87]). It is this unconscious substratum that is the seat of latent beliefs and emotions that, on being activated, make possible the inferential activities within the mind of the individual. The processes of acceptance, testing, and development of scientific theories by a scientific community in any specific area of investigation also has a deliberative component wherein a theory is thoroughly criticized, analyzed, and tested against available evidence before acceptance as a valid theory. However, this entire process is conditioned by *underlying* beliefs and value judgments located in a cognitive *substratum*, where the substratum is made up precisely of the *individuals* of the community who carry within their minds the latent beliefs and values of the society at large. Analogous to the (computationally) parallel processing within the cognitive unconscious of an individual, one has an essentially similar process of exchange and communication among individuals and groups of a scientific community, based on which the deliberative activity of the community as a whole finally takes shape.

Ronald N. Giere, in his book [43], draws attention to ‘collective cognition’, or *distributed cognition* in modern research establishments such as the European organization CERN where a large number of individuals, along with computers and various instruments making up an enormous set-up, form a huge cognitive system. The individuals operate, to a large extent, in parallel, communicating with one another by means of language. In reality, cognition may be distributed over wider networks such as a scientific community or even a more widely spread society sharing a common culture.

The dominant conditioning effect on the scientific deliberations of a community in some specific area of investigation is exercised by means of a *paradigm* ([76], chapter 5) which is in the nature of a tacit *way of approach* or a *point of view* in the problem area under consideration. Put differently, a paradigm is a heuristic of major relevance, or an *exemplar*, a tacitly accepted conceptual kernel, that acts like an organizing principle around which various conceptual elements get crystallized into a coherent whole, thereby giving rise to a complete theory and its ramifications. While the paradigm is a *conditioning* factor, the major *active* factor in the formation of a theory is, of course, a hypothesis that initiates the process of theory building and itself acts as the conceptual focus of the theory. This distinction between a conditioning factor and an active component in a cognitive process is a useful and relevant one.

Strictly speaking, a major hypothesis or a new theory emerges in violation of a reigning paradigm whereby a new paradigm is brought into existence. Still, the distinction between a conditioning factor and an active factor in inference making and theory building is one of relevance.

Kuhn introduced the idea of a paradigm in a specific sense, namely that of an exemplar of major significance.

“Close historical investigation of a given specialty at a given time discloses a set of recurrent and quasi-standard illustrations of various theories in their conceptual, observational, and instrumental applications. These are the communitys paradigms,

revealed in its textbooks, lectures, and laboratory exercises. By studying them and by practicing with them, the members of the corresponding community learn their trade” ([76], p 43).

However, it is of considerable value to extend the idea to one of a *tacitly held guiding principle* in the broader context of cognitive science and in the study of scientific practice. Kuhn himself underscored the tacit aspect of a paradigm by asserting that the latter was not a set of principles that could be spelt out in explicit logical terms. In any case, the idea of a paradigm is so fertile that it has been stretched, interpreted, and adapted in a great many ways and in a great many contexts, though always around a core meaning, namely, the one of a tacitly entertained exemplar. Kuhn himself used the idea in subtly different ways and I plead guilty to having taken similar, if less deserved, liberties in the present book. However, in all such usage, I have tried to specify the meaning by implication to the extent possible.

While the paradigm is a major conditioning factor in the process of theory building by a scientific community, other latent factors also play a role, where these relate to *cultural* influences of the society at large and the social *value system*. All these latent effects operate through the cognitive substratum of the community provided by individuals, institutions, dominant groups, and *power* structures. Taken together, these make up the *mode of thinking and reasoning* of a scientific community and set the *goal* of scientific exploration in various areas of inquiry.

The aim and method of science: the question of values

The question often arises as to whether science is a *rational* enterprise. Rationality of an activity or an approach involves two considerations — first, a specification of the *aim* of the activity, and secondly, an understanding of the *method* followed in the activity with reference to the aim, i.e., in other words, an appraisal of whether the method is appropriate with reference to the aim. The commonly assumed rationale of mankind's scientific activity is often stated in terms analogous to the way one evaluates individual inferential and decision-making processes: science, it is said, is motivated by

a combination of *instrumental* and *epistemic* aims. Instrumental rationality is motivated by the aim of being *effective*, where the effectiveness is judged with reference to whatever *goals* may set the activity — either individual or collective — in motion. Epistemic aim, on the other hand, relates to the acquisition or use of *truth*. For instance, investigations on the toxic effect produced by a certain bacterial specie may be undertaken with the aim of understanding ways to treat a disease caused by it, and the effectiveness of the methods followed in these investigations may be judged by checking whether a means of treating the disease is actually realized in the process. On the other hand, the epistemic rationality of the investigations is to be judged by checking what new pieces of knowledge relating to the life processes of the bacterium and to the cellular and biochemical causes linking the bacterial toxin with the disease in question are gained in the process. Generally speaking, epistemic virtues reside in an instrumentally effective inferential and decision-making activity, since effectiveness is rarely achieved without a good understanding of the relevant part of reality. The commonly cited goal in the case of science is ‘human welfare’, in the service of which science sets about gaining knowledge, or a true understanding, of the world. At times, the epistemic aspect is cited as being the fundamental criterion of doing science, and philosophers and workers in foundational issues in science are fond of underlining its truth-seeking aim.

In reality, the goals and motivations underlying the scientific enterprise in the modern world are complex, conflicting and, at times, dark. There is the individual scientist, there is her superior, there is the institutional head, there is the fund-giving agency, there are governmental policies, there are dominant interest groups, there is public opinion, there are media campaigns, and there are factors operating on an international scale. In every area of human interactions there are, moreover, *power relations*. Goals of scientific investigations in various areas of interest are ultimately set by this entire complex of factors, and often appear in contrary forms, depending on who is viewing the goals. The individual scientist may look at her research project as a means to gain genuine understanding of the mechanisms underlying some natural phenomenon, which means that her personal goals are probably epistemic in nature, though these need not be *purely* epistemic since there may be admixtures of career ambitions, monetary as-

pirations, peer rivalry, craving for peer recognition, and similar other 'worldly' motives in her scientific pursuits. The institutional head and the fund-giving authorities, on the other hand, may have completely different perspectives, not unconnected with interests of dominant power groups. This brings us to the question of *values* underlying the business of science. We will return to this later in this essay and will see that there is no *logical* or *preferred* approach to this question as there is no such approach to the question of rationality of science and of scientific *method*. Questions relating to science as an individual and social endeavor are as open as those in any other field of human concern. One can never sit back and engage in the consultation of charts and graphs to pronounce on the correctness or otherwise of the course of science.

Interestingly, the *epistemic* concern continues to remain as one of constant relevance among all these conflicting pulls and pushes on science. *One has to correctly understand the atomic nucleus in order to drop a bomb*. Even in its most worldly and dark pursuits, science has to go for truth. And truth, as we have seen, rests on *interpretation* — the process of basic relevance when cognition confronts the world. On the one hand, truth has to aim for a correspondence with elements of reality while, on the other, it has to build upon prior *experience* of reality, a prior conceptual structure, and prior beliefs and preferences resulting from individual and social inferential processes. It is a disproportionate concern with the second, interpretational, aspect of truth that produces the idea that truth is *constructed*. Once again we fall prey to a fundamentally dichotomous way of thinking resulting from our craving for *simplicity*. In order to avoid the complexity and strain of examining an issue having conflicting and contrary aspects to it, the mind tries to grasp the *dominant* aspect. But the 'dominant aspect' is somewhat a myth — it depends on the point of view.

Science as a telescope: the 'objective' and the 'eyepiece'

The point of view of realism tends to grasp the reality-oriented aspect of truth while the point of view of anti-realism harps on the cognitive and constructive aspect. An illuminating analogy would be to think of the 'objective' and the 'eye-piece' of an astronomical

telescope. In viewing a remote object, the objective lens system of a telescope is aimed at the object, while the eyepiece lens system focuses the rays on to the eye. The telescope itself cannot be reduced either to the objective or to the eyepiece. I will now consider in succession the ‘objective’ and the ‘eye-piece’ of science — its truth-seeking aspect and its interpretational aspect, keeping in mind the admonition that science is neither the one nor the other, but a complex whole made up of the two.

The reality-oriented aspect of truth is, moreover, *contextual*. The truth we learn about nature is always partial, with ever-expanding scope. Imagine the objective of a telescope — let’s say, of one of Galileo’s favorites — aimed at various different celestial bodies, while the magnification and the resolution achieved by means of the instrument get increased from time to time. In the process, there emerge ever newer pictures of the firmament, spectacularly causing kaleidoscopic changes in our conception of the heavens. This, then, can be taken as a metaphor (and, at the same time, an *instance*) of how science acts as a telescope designed to probe nature, with its objective constituting the external context of observations and theory building, seeking to arrive at ‘truth’.

Continuing to refer to the actual telescope (an imagined one, though), the eyepiece handles the job of collecting the rays passed by the objective and focusing those on to the eye and, in the process, it introduces *aberrations* and *diffraction* effects, whereby the image formed by the telescope differs from an exact and ideal representation of the object and only *corresponds* to it in some complex manner. In terms of the *allegorical* telescope representing the probing of nature by science, the aberrations and diffraction effects that depend on the detailed structure of the eyepiece are indicative of the role of our existing conceptual framework — the *internal* context of observations and theory building — in the process of making up the interpretation of nature that science provides us with.

As our concepts and our instruments capture newer and newer aspects of nature, scientific theories are said to *approximate* more and more an ideal and perfect description of reality ‘out there’. However, the exact sense in which science *progresses* in its endeavor to unearth the mechanisms of nature has raised controversy, and some strife.

Approximations in science: the metaphor of convergent and divergent series

Imagine a sequence of numbers approximating the irrational number 'pi', defined as the ratio of the circumference and the diameter of a circle. The rational number 3.14 is a good approximation for many purposes, but a better approximation is provided by the number 3.1416, where two more decimal places have been added. Even this, however, does not constitute an accurate description of 'pi', and the number 3.14159 constitutes a still better approximation. The quest for capturing the exact value of 'pi' does never end, since the successive rational approximations never suffice to produce the exact value which, in the language of mathematics, is an irrational one. What one needs is to set up a *convergent series* for 'pi' and go on adding up to successive terms of the series, numerous such convergent series being known for approximating the elusive number. Interestingly, the number, while being elusive, has an indubitable existence as an irrational number (and a transcendental number at that), it being only the case that any attempt at capturing its value *in terms of a rational number* can never meet with final success.

Does science approach a description of reality in successive waves of cognitive advancement in an analogous manner? It is often said that science is approaching more and more towards an ultimate truth about nature, though that truth will forever remain beyond our grasp. Is the progression of science, then, something analogous to an approximation of 'pi' by means of a convergent series? In reality, the progression is of a much more complex nature, somewhat analogous to the approximation of a function by an *asymptotic series*.

Asymptotic series and singular approximations

Consider a function of one single variable like, for instance, the distance (s) covered by a car as a function of time (t) where the latter is the independent variable and the former the dependent one. A complete knowledge of the function would enable us to

depict the relation between the two by means of a graph. But suppose that such a complete knowledge is not given to us and that we attempt an approximation by means of a power series, aiming at a good description of the graph at all points close to some pre-assigned value of t , say, near $t = 0$. Now, the power series is in the form of a sum of an infinite number of terms and one can, for any given value of t (close to $t = 0$), add up the successive terms, hoping to approach the corresponding value of the dependent variable s in a manner analogous to the way the value of 'pi' can be approached by means of a convergent series.

An *asymptotic series* does provide an approximation, but one of a more *complex* nature. As one goes on adding successive terms for any given small value of t , the series at first appears to converge to some value, but when a good number of terms have been added, the sum begins to *diverge*. An excellent approximation to s for some given value of t can be obtained by summing up, say, the first ten terms of the series, but the quality of the approximation *deteriorates* as one sums up the subsequent terms while, for a larger value of t , a larger number of terms may be needed in order to obtain a reasonably good approximation.

Scientific investigations often aim at guessing the relations between dependent variables and relevant sets of independent variables (there may be more than one independent variables in any given problem) from observed data, and any such relation arrived at by analysis of the data is typically expressed in the form of a series. However, a relation depends, in general, on a set of relevant *parameters* that are indicative of the *conditions* under which the proposed relation holds. One can now ask the question as to how the relation under consideration depends on the value of some parameter, which we denote by the symbol p ? This is of a more involved nature than the one of the graphical relation between a dependent and an independent variable, since now the graph itself gets changed as the parameter varies. If the parameter is made to change by a small amount, the graph also changes, which is commonly expressed by saying that the relation between the dependent and the independent variable gets *perturbed* to a small degree by the variation of the parameter. One can again express the perturbation in the form an approximation by means of a *power series* where now the power series contains

terms with successively higher powers of the parameter p .

In numerous instances of practical interest, the power series turns out to be *asymptotic* in nature, in which case one is left with a *singular* perturbation around some reference value of p , say, $p = 0$. For instance, when a plane wave of light, corresponding to a bunch of parallel rays, is made to pass through a circular aperture and then captured on a screen placed on the other side of the aperture, a distinctive pattern of variation of light intensity (dependent variable) with distance (independent variable) along any line passing through the center of the illuminated patch is obtained, where the graph of the variation is found to be *oscillatory* in nature, the oscillations being confined to a narrow region near the border of the illuminated patch. The relevant parameter here is the ratio $p = \frac{\lambda}{a}$, where λ stands for the wavelength of light used and a for the radius of the circular aperture. The oscillations are indicative of the *wave nature* of light, due to which the parameter p typically has a small but non-zero value. On the other hand, the limit $p = 0$ corresponds to an *idealization* in the description of optical phenomena that goes by the name of ‘ray optics’ or ‘geometrical optics’. In this ray optic description, where the wave nature of light is ignored, the variation of light intensity is depicted by a uniform intensity inside the illuminated patch, which abruptly drops off to zero value at the border of illumination, with the zero intensity continuing into the shadow region on the screen. This differs *qualitatively* from the oscillatory variation that obtains for any non-zero value of p , *however small*. This is a well known singular perturbation problem in physics. It tells us that the description of optical phenomena in terms of the wave theory is *fundamentally* different from a description in terms of the ray theory, because the latter ignores an *essential* aspect of light.

In an optical experiment of the type indicated above, the ratio p sets the *context*. In a large class of experiments the actual value of p is so small that one can assume p to be zero and still come up with a good number of predictions arrived at with the help of the ray theory, but all such predictions fail when one focuses on a narrow region around the border of the illuminated circular patch, where one observes oscillatory variations in light intensity (it is in this sense that one can say that the ray theory, while being a useful and effective one within a certain context, enjoys only a *skewed*

fit with reality). In order to investigate these oscillatory variations, one has to design experiments with apertures of small radius (in which case the non-zero value of p can no longer be ignored), which corresponds to observations in a different context — a context where a new theory is needed *beyond* ray optics and that theory is commonly referred to as ‘wave optics’.

This is how the question of *context-dependence* in the description of reality by scientific theories comes up. A theory that suffices in a given context is to be replaced with one of a different *foundational structure* when the context changes in such a way that previously ignored aspects can now be ignored no longer, and a consideration of those aspects makes necessary a modified theory with an altogether new texture. However, the older theory remains *woven into* the newer one in a relation of some complexity, as the ray theory is woven into the wave theory of light, where the complexity often finds expression in terms of an asymptotic approximation.

Singular reduction from one theory to another

A much discussed case of such singular transition from one theory to a succeeding one is that from the *classical* to the *quantum mechanical* description. In order to illustrate the singular nature of the relation between the two types of descriptions, it is illuminating to refer to a particle moving along a straight line and encountering a *potential barrier* — a region where the classical theory forbids the particle from entering because of insufficient energy possessed by it. This stricture of classical mechanics is, however, *conditional* on the *Planck constant* (h) being ignored. However, the Planck constant, though having a small value, *cannot be ignored in all contexts*, and there do exist situations where a particle, on encountering a potential barrier, can be seen to *tunnel* through it with a small but non-zero probability (the process of emission of alpha particles from nuclei occurs by means of such quantum mechanical tunneling). The dependence of the probability of tunneling on the parameter h turns out to be of a singular nature, where the behavior of the particle for a small non-zero value of h differs qualitatively from its behavior when h is ignored as not being of any consequence. A similar story, indicative

of a singular relation between the classical and quantum theories, can be told regarding the celebrated formula for the *black body spectrum* first derived by Planck.

The quantum context is thus distinguished from the classical by the non-zero value of the Planck constant h . In any context in which the Planck constant is negligibly small compared to the relevant *action(s)* (a physical quantity characterizing a dynamical system, analogous to its energy or momentum), a classical description suffices (though, only conditionally) in describing the behavior of the system, while in an altered context, where h can no longer be ignored (in the domain of microscopic systems, for instance) a different description is necessary where, from the mathematical point of view, the transition from the former description to the latter is singular in nature.

Not all transitions from a more restrictive theory to a less restrictive one (the classical theory is more restrictive as compared with the quantum theory because it holds in a restricted context where the Planck constant can be ignored as being negligibly small) are however, of a singular nature. For example, the transition from Newtonian mechanics to the mechanics based on the special theory of relativity is, from the mathematical point of view, non-singular, and special relativistic formulas in mechanics go over to Newtonian formulas in the limit of relevant velocity ratios ($p = \frac{v}{c}$, where v denotes any relevant velocity characterizing a dynamical system of interest, and c stands for the *velocity of light in vacuum*) going to zero, in a smooth or non-singular manner. In other words, the behavior of a dynamical system as described by the special theory of relativity for small but non-zero values of $\frac{v}{c}$ does not differ qualitatively from its behavior as described in the Newtonian theory. Still, the special theory of relativity involves a major conceptual reorganization as compared to the Newtonian theory.

Indeed, a major conceptual restructuring becomes necessary whenever some relevant aspect of reality, which was ignored in the earlier theory, now assumes significance, in consequence of which, parts of reality appear *in a new light*, as if a new *dimension* is added to reality. Often, the new 'dimension' becomes 'visible' in virtue of the fact that some specific natural parameter (p), which was not taken into account in the old theory because of a restricted domain of observations, now assumes an essential role in an

expanded domain. This is the case in optics where the ray theoretic description fails in a domain in which the finite non-zero value of the ratio $p = \frac{\lambda}{a}$ (here a stands for a typical length scale in the observational context) becomes significant. In quantum theory, on the other hand, it is the non-zero value of the ratio $p = \frac{h}{A}$ (A stands for the typical value of the action of the dynamical system under consideration) that assumes relevance. And, in the case of the special theory of relativity, the ratio $p = \frac{v}{c}$ plays the corresponding role. In each of these cases, an attempt to force nature into a straight-jacket by assuming that p is negligibly small, results in anomalies and inconsistencies whenever the context of observation needs the non-zero value of p to be taken into account. In other words, there are certain *natural borders* that *cannot be crossed* (from zero to a non-zero value of p) without an appropriate conceptual restructuring, since a new ‘dimension’ or *aspect* of nature is revealed across any such border: nature now appears in a new *perspective*.

The idea of *perspectival realism* espoused by Ronald Giere in [43] is analogous to that underlying the context-dependence of scientific theories.

The crossing of borders: scientific revolutions

Events of restructuring of theories resulting from an alteration in the observational context and the crossing of such natural borders in various different fields of scientific investigation are momentous ones. One feels inclined to identify these with the scientific *revolutions* that Thomas Kuhn spoke of ([76], chapters 9-13). There have been numerous protestations from philosophers of science questioning the validity of the concept of scientific revolutions, but many of these are in the nature of *logical* and *analytical* discourse where events in real life, belonging to the real world, are viewed in the abstract, and stringent demands are imposed on the way these are to be analyzed and classified. Nothing in this world can be described in terms of *pure* categories, devoid of conflicting and contrary aspects, and every description is valid only contextually, and the description of events of conceptual restructuring in science as revolutions is no exception. For one thing, revolutions occur in *all scales*. In order to understand this, one is to look at periods of scientific activity that Kuhn referred to as *normal science*,

in contrast to revolutions. As with the idea of scientific revolutions, Kuhns reference to 'normal science' has also raised controversies because of a tendency to assign undue and disproportionate significance to the concept of normal science. Normal science refers to much of the day-to-day activity of scientists where conceptual restructuring is not needed and where one can discern an unspoken allegiance to some conceptual *paradigm*. All these ideas of paradigms, normal science, and scientific revolutions are nothing more than descriptions, in the spirit of naturalism, of the way things appear to be in the practice of science, in respect of its epistemic aspect. These tell us how things *are*, and are not proclamations of how things *ought to be*, though the distinction between the two, which is not always a sharp one, seems to have been disregarded, at times, by Kuhn himself, perhaps by oversight rather than by design. In particular, normal science is *no subservience to the authority* of people who have been responsible, fortuitously or otherwise, for the introduction of new ideas and concepts in a field.

Normal science: conceptual restructuring at all scales

In real life, one often follows a certain set conceptual pattern till a new and more versatile pattern makes its appearance. We often benefit from following a set pattern, though the danger of falling a prey to the pattern is, of course, always there. Indeed, seeds of new concepts accumulate in an invisible process even as one is engaged in the mundane practice of following a set conceptual pattern, and what one needs is to recognize these as they germinate so as to make possible a conceptual reorganization. In other words, periods of following a set conceptual pattern and events of conceptual reorganization are inextricably linked with one another, and a distinction between the two can be made only provisionally and within given contexts.

Indeed, quiescent states in which seeds of new ideas are nurtured, and phases of conceptual restructuring, alternate *in various scales*. What appears to be a period of normal science from a bigger perspective, involves definite conceptual changes when looked at more closely, where, however, concepts undergo a revision on a smaller scale, within a larger framework that remains unchanged. For instance, within the framework of the

basic ideas of *fluid dynamics*, there have been vigorous attempts at introducing new concepts for the understanding of the remarkable phenomenon of *turbulence*, and a number of new ideas have indeed been introduced that have had far-reaching consequences even outside the field of fluid dynamics. Research on turbulence has progressed in the spirit of normal science within the framework of basic ideas in fluid dynamics (where, in a manner of speaking, the paradigm is set by what is referred to as the *Navier-Stokes equations*), but when looked at more closely, there have taken place conceptual revolutions within the area of turbulence, a notable example of which is constituted by a number of fundamental ideas of Kolmogorov and others introduced some seventy years back.

Incidentally, the concept and study of turbulence in fluid dynamics is also associated with a certain parameter of crucial relevance, namely the *Reynolds number* (or, in a different context in fluid flow, the *Rayleigh number*): as a critical value of the parameter is crossed, the nature of the flow changes qualitatively and, from a mathematical point of view, *singularly*. Once again, the Reynolds number can be said to set the context in which one investigates the mechanisms and characteristics of fluid flow.

The complex relation between theories: emergent phenomena

As theories succeed one another in any given field of investigation, consequent to contextual changes relating to the domain of observation of natural phenomena (such as, observations for small and large values, respectively, of the parameter $\frac{\lambda}{a}$ in optical set-ups), there occurs a conceptual restructuring where two consecutive conceptual frameworks do not bear a relation of direct correspondence. Of the two theories in question, the earlier one usually applies to a more restrictive context (in a relative sense) as compared to the succeeding theory, whose concepts are of a broader scope. It is often the case that many of the concepts of the earlier theoretical framework do not have counterparts in the later theory, though those can be interpreted within the new and expanded theoretical framework. It is often said that the former theory is reduced by the latter

(i.e., the later one), or that the latter theory reduces to the former (for instance, in the case of succession of the classical theory by quantum theory, by restricting the description of phenomena to a domain where the ratio $\frac{h}{A}$, introduced above, can be ignored). However, the idea of ‘reduction’ often involves considerations of a complex nature, and is associated with the concept of *emergent phenomena* (such as the phenomenon of tunneling that emerges across the border separating classical and quantum mechanical descriptions), commonly associated, from the mathematical point of view, with *singularities*.

Michael V. Berry, the great mathematical physicist, spoke of asymptotics, singularities, and theory reduction in a remarkably suggestive cameo article in his typically lucid and insightful style.

“In science we strive to integrate our experiences, observations, and experiments into a single explanatory framework — ‘a theory of everything’. Of course this goal has not been achieved, and probably never will be. What we have instead are the partial descriptions provided by biology, chemistry, physics, etc., and, within these, the various sub-fields such as fluid mechanics and quantum mechanics. The different areas of study do not fit tidily together. Particular difficulties arise when a more general description is supposed to encompass an older, less general, one, usually by providing a microscopic explanation of its principles. It is hoped that a less general theory can thus be ‘reduced’ to a more general one. But this comfortable picture is often spoilt by certain classes of higher-level, or ‘emergent’, phenomena..... .”, [9].

Berry speaks of the relation between two theories, of which one is supposed to get reduced to the other in the limit of some parameter (δ) going to zero (analogous to the parameter p introduced earlier) where, however, this limiting transition is often not a *smooth* one.

“We shall see that very often reduction is obstructed by the fact that the limit is highly singular. Moreover, the type of singularity is important, and the singularities are not only directly connected to the existence of emergent phenomena but underlie some of the most difficult and intensively-studied problems in physics today.”

In this book I have tried to adopt an approach in keeping with the spirit of the ideas underlined by Berry in this article of his, which is likely to prove to be one of major importance in the philosophy of science. He considers a number of concrete examples of the relation of reduction between theories in physics, where the ‘reduction’ is more complex than what at first appears to be the case.

It is the complex relation between prior and succeeding theories in various fields of scientific investigation that can be subsumed under the idea of *incommensurability* of theoretical frameworks, mooted by Kuhn.

“The normal-scientific tradition that emerges from a scientific revolution is not only incompatible but often actually incommensurable with that which has gone before.” ([76], p 103).

Even as two theories appear to be incommensurable with reference to each other, there still remains an enveloping language by means of which the two theories can be examined side by side.

Feyerabend also had much to say on the notion of incommensurability, though not quite in the same context ([38], chapter 15). Polanyi spoke of incommensurability in his own terms of reference, and his ideas may have had influenced Kuhn and Feyerabend in the formulation of their views (see [66]).

As with the other ideas that Kuhn introduced, the one of incommensurability has been the focus of quite considerable controversy (reviewed in [111], [32]). But, once again, this idea of Kuhn's finds a natural —it resonance with what one *feels* as one goes through histories of theories succeeding one another in the annals of science. It is precisely because of this resonance that Kuhn was able to strike in the minds of men, at a level deeper than that addressed in much of formal, analytical, and philosophical discourse, that his work ushered in a new era in the study of science as such. Of course, Kuhn did no more than point at the basic idea and did not make it much more explicit, and the idea itself may get transformed in days to come, but the fact will remain that

Kuhn did point at something of central relevance in the study of how scientific theories appearing in succession are related to one another. A number of concepts likely to be of considerable significance in such a study are the ones relating to contextuality, reduction, singularity, and emergence (refer to [6], [19], [13]).

The notion of contextuality (as mentioned earlier, this term is used to mean context-dependence) of scientific theories and of contextual truth that I have tried to outline in the present essay can now be summed up.

Contextuality in science: summing up

In any sphere of scientific investigation and at any stage of evolution of scientific ideas, some part of reality is studied within some definite context that is, typically, not recognized explicitly by the investigators, since there remain transitional borders that go unnoticed. With an expansion of the domain of observation of facts and phenomena, and with the gradual evolution of ideas within a given theoretical framework, an altered context emerges where previously unrecognized borders are recognized and crossed, new aspects or 'dimensions' of nature are revealed, a new perspective emerges, and the texture of the relevant theory gets transformed. There results an *altered conceptual structure* in which many of the concepts of the previous theory are replaced with new ones, though these former concepts can be interpreted from the vantage point of the newly emerging theory. The earlier theory is often subsumed within or reduced by the succeeding one where singularities and emergent phenomena are commonly met with. In a sense, the earlier and the later theories are incommensurate with each other. Though incommensurate, both are of epistemic value within their respective contexts. The truth of either theory is contextual but not relative in the sense of being principally a matter of arbitrary interpretation by individuals or communities.

This last notion of truth as being relative to interpretation or construction, and not as one of referring to reality is commonly associated with the anti-realist view of science. However, we have seen earlier in this book that, even from the realist point of view, while truth is focused on to reality, it has a second facet, since it emerges in a protracted

cognitive endeavor, which is basically in the nature of an interpretation of reality.

This admixture of the element of interpretation in scientific theories makes the latter suspect when judged against the criterion of *objectivity*, but that, of course, is no reason to simplify the notion of truth by ignoring its interpretational aspect. As we have seen, truth is a complex thing where its correspondence to reality is conditioned or *constrained by* its cognitive origin, where cognition is nothing more than an interpretation while, conversely, it is an interpretation *aimed at* and *answerable to* reality. This, as far as I understand, is the naturalist perception of truth — truth as it is, not truth as it *ought to be*.

The naturalist point of view examines and addresses the idea of scientific truth by recognizing both its cognitive-interpretative aspect and its commitment to reality, which entails complexities and tensions in the notion of truth. Philip Kitcher recognizes the two contrary aspects to truth and has examined these from various angles. Based on Kitcher's ideas on the notion of scientific truth, one can have a broad view of the issues involved in an authentic discourse on scientific theories as bearers of truth, though that discourse may eventually turn out to develop approaches and interpretations differing somewhat from Kitcher's. I have found [49] to be a useful source-book for this.

Theory choice: the problem of 'underdetermination'

This is supposed to raise a problem with *theory choice*. How do we choose a theory as being a correct one, especially as we recognize that, when judged against available evidence in any given domain of scientific investigation, there may be not one but, in principle, an infinite number of theories compatible with the evidence (the problem of 'under-determination' of theories by evidence). In reality, it is *seldom* the case when the scientific community faces the task of *choosing between alternative theories*, all compatible with the body of observed facts in some given domain of inquiry, though alternative theories can conceivably be constructed. However, it often happens that several alternative *hypotheses* are offered to explain *anomalies* in the area under consideration,

one of which survives criticisms and scrutiny from various quarters and, in some special cases, a number of mutually compatible hypotheses that cohere with one another evolve into a theory. A hypothesis is a result of an act of abduction, and abduction by its very nature is a guesswork aimed at producing a sufficiently consistent and effective explanation of a number of facts of observation. It is an act of cognition which, in turn, is based on interpretation against a body of items of knowledge and beliefs accumulated in past experience. There may, in principle, be numerous possible hypotheses consistent with facts of observation, but only few of those are *also* consistent with the knowledge base and the beliefs that make up the internal context of the cognitive act, and fewer still survive the intricate and subtle *process* of hypothesis generation, some of whose salient features have been outlined in chapter 8. Thus, a hypothesis is not an exercise in logic alone, but a complex and intricate inferential process constrained by a great many cognitive factors conditioned in a long history of evolutionary development.

All the crows I have observed in the city of Kolkata during my lifetime have been black. From this I generalize that all crows are black. I could also generalize to the effect that all crows in the city of Kolkata are black, and those in the city of London are white, which is *logically* compatible with the facts of observation on which I generalize. Or I could choose from an infinity of similar 'generalizations', but the latter do not conform to features of past generalizations that have been found to be fruitful and effective (I do not enter into the question of the biological features that define a crow and assume that a crow is unambiguously identified by sight, without reference to its color), and so I do not entertain these. The cognitive apparatus within me has developed certain processes of arriving at *relevant* and meaningful generalizations (meaningful not in the sense of logic but in the sense of making sense of this world and surviving in it), which make most of the above generalizations irrelevant in the context of my inference.

Imagine for a moment the phase of history when the framework of classical physics was being put together, principally around Newtonian mechanics and electromagnetic theory. All the observational facts collected up to that point of time were *compatible* with the theory of relativity and with quantum theory as well. Why, then, were the theory of relativity and quantum theory (or, to indulge in a further flight of fancy, the

quantum field theory) not considered and adopted by physicists and mathematicians of the time who were equipped with astounding reasoning ability? The reason is precisely the same as that in the case of your child suddenly crying out in an adjacent room, and you making the inference that she must have fallen and hurt herself while *not even imagining* the possibility that a snake may have entered the room and bitten her which, however, is *compatible* with the fact of the child having cried out. In either case only those theories and hypotheses were entertained that, in addition to being logically possible, were consistent with the context. Hypotheses are made *in accordance with the context*, building upon ideas and notions that have accumulated in past experience (while, at the same time, bringing in novel elements as necessary) with the aim of solving *specific anomalies*, puzzles, and problems, because that is the way that human cognition works. In making a hypothesis, the cognitive mind does not think of anomalies and problems that may come up in a different context — one that is not of current relevance. The black body problem, or the puzzle of photo-electricity, or the frame-independence of the velocity of light were not of relevance in the context of eighteenth and nineteenth century physics, which is why the quantum theory or the theory of relativity were outside the realm of possibility at that stage even though these were all compatible with facts of observation accumulated up to that time — no fact of observation were logically in conflict with these (and, what is more, these subsequently turned out to be successful theories, contrary to the example of the crows mentioned above).

Hypotheses and theories are not plucked out of thin air but are built in steps upon past successes in much the same way that natural selection acts in the emergence of new species — bringing in novelties within the existing pool of genetic material.

Looked at this way, a scientific hypothesis is hugely constrained by the *internal* context of knowledge and beliefs of an individual and of the society, while being consistent with the external context set by facts of observation, where certain aspects of phenomena go unnoticed. Consequently, only a few distinct hypotheses can, in practice, compete with one another in any given field of investigation at any given point of time. As with the beliefs of an individual, numerous different reasoning attitudes, modes of thought, and half-formed scientific ideas, along with relatively remote beliefs, set the internal context

of a scientific community in respect of its choice among such competing hypotheses. All these factors making up the internal context play a role in defining the attitudes and preferences of various different sections of the scientific community towards the competing hypotheses, as a result of which the choice among these hypotheses assumes the form of a complex and often protracted process but, in the end, the one of greater epistemic value usually prevails, and a new theory is built up around it.

While a scientific theory has to be, in some sense, a true description of some aspect of nature (recall, however, the complexities inherent in the concept of truth of a theory), a more immediate requirement that it has to satisfy is that of *consistency* — consistency with the existing body of knowledge and with observational and experimental evidence relating to the relevant field of inquiry. Theories are ultimately accepted or rejected on grounds of consistency, though a consistent theory does not necessarily attain a desired standard of truth. In contrast, an emerging *hypothesis* is often not consistent with the existing body of knowledge, since it bears the stamp of a belief. It is accepted by a section of working scientists only because it *promises* to reveal novel ‘dimensions’ of nature. If the hypothesis does succeed in delivering the goods, then it finds its place in some *new* theoretical framework, where it meets with the requirement of consistency within an emergent theoretical framework.

The new theory does not appear in one single package ready for the market. Once a number of anomalies appear within an existing theoretical framework (*explanation of the black body spectrum, explanation of the line spectrum of hydrogen*), there occur many failed attempts at framing a hypothesis that can explain at least some of these anomalies. In the case that more than one hypotheses are offered, the question of choice among these comes up, which sets in motion currents of criticism, appraisal, and analysis, where non-epistemic factors play their role alongside epistemic ones, as indicated above. At times, several hypotheses are offered and found useful (*Plancks hypothesis, Bohrs hypothesis*) before these are collated to form the nucleus of an emerging *theory* (*quantum mechanics*).

In this respect we once again observe a similarity in the ways in which a hypothesis is formed in the mind of an individual, and a theory takes shape in the scientific practice of a scientific community. Each of the two *processes* involves factors of a non-epistemic nature, based on prior attitudes, emotion-linked preferences and beliefs, but each, at the same time, involve epistemic judgments and consistency checks, as a result of which the *end product*, in the form of a hypothesis or a theory, has to be of epistemic value (while being, at the same time, *conditioned* by non-epistemic factors). As we have seen, this epistemic value is not generally in the nature of a *logical* soundness, but one relative to current aims and goals and relative to the *context* (such as the one set by the Reynolds number in fluid flow, or by the parameter $\frac{\lambda}{a}$ in an optical set-up) in which some part of nature is observed and studied.

Scientific progress: socially determined or socially conditioned?

The anti-realists put emphasis upon the non-epistemic aspects of theory choice so as to assert that the succession of theories in the evolution of scientific practice is *socially determined* — a point of view apparently conforming to the views of the likes of Kuhn and Feyerabend [76], [38]. In particular, Kuhn's observation that successive theories are incommensurate between them and that the progression of scientific theories is not of a cumulative nature converging to some ultimate truth are, *on the face of it*, supportive of this anti-realist point of view. However, we have seen that these notions do have a *realist* interpretation as well, where the incommensurability and the non-cumulative nature of scientific progression are consequences of the contextuality of theories in the sense outlined above. Science has a progressive aspect to it since it unravels the mechanisms of nature layer by layer, across a succession of natural borders, and the metaphor of a convergent approximation to some final truth has to give way to metaphors of a more complex nature such as the one of approximation by means of a succession of asymptotic series — each describing a limited approximation in some specific context, where it is conceivable that contexts appear in an unending succession: Nature has an infinite number of distinct 'dimensions', and is inscrutable to the end.

The question as to whether and how one can locate continuing progress in the history of science has been a contentious one. In particular, Kuhn and Feyerabend have been interpreted as having been opposed to the idea that there has been a continuing progression or development in scientific theories and practices.

In the case of Feyerabend, such an interpretation of his views has been fostered by his contrary way of saying things — a way that was meant to jolt people into questioning the received view. In fact, he had too much of sense in him to question the continuing discoveries in science in a naive way. What he did want to highlight is that there cannot be a predetermined yardstick of progress in line with the commonly perceived and advertised ‘methods’ of science. He, above all, was against the mindlessness in the name of ‘intellect’ in science.

“ ‘Progress of knowledge’ in many places meant killing of minds.” ([38], p 3).

Kuhn, on the other hand, was remarkably lucid in what he meant to say, but the import of what he said was, at times, not transparent because of its novelty. Moreover, he was too engrossed in bringing in new ideas in bold strokes of the brush to finish these with detailed explanations. He was fully cognizant of the *continuing evolution* in science, and never questioned the fact of scientific progress but, at the same time he was concerned with setting right the underlying notions. He advanced the insight that the scientific process achieves distinction by way of eliminating, in a large measure, disagreement on *foundational* questions. Generally speaking, progress depends crucially on creative contributions, and scientists in any given area of research are capable of identifying and agreeing upon these creative contributions precisely because they share a common paradigm and are generally reluctant to engage in controversies upon foundational issues. It is only in the backdrop of normal science that scientific revolutions make sense, and progress in science is tied with the occurrence of such revolutions. What is novel in Kuhns discourse on progress is that progress occurs through a succession of theories *incommensurable* with reference to one another, and is therefore of a *complex* nature ([76], chapter 13) where, moreover, the complexity repeats itself at various scales nested within each other.

The spectrum of views on scientific progress, to be found in the philosophy of science

literature is discussed in [90].

In other words, it is possible to recognize and accept the *fact* of scientific progress (or of a continuing evolution and development in the sciences) while at the same time raising deeper questions as to the *nature and relevance* of that progress. These deeper questions, however, need not imply that the evolution of scientific theories is driven predominantly by *non-epistemic* factors — ones without valid reference to mechanisms inherent in nature.

But, like an albatross, the non-epistemic factors involved in the interpretational aspect of truth cannot be shaken off, which is why the traditional dichotomy between realism and anti-realism is not considered to be of overriding relevance in the naturalist approach to the philosophy of science. The naturalist point of view adopts the approach of science itself— when confronted with a choice between apparently conflicting descriptions of reality, science delves deeper and comes up with a broader, if more complex, description. The case of apparently contrary aspects — epistemic and non-epistemic — in the notion of truth is of a similar nature, needing a deeper and broader view.

The naturalist approach acknowledges the involvement of individual and social beliefs, cognitive preferences, and modes of thought in the interpretational aspect of scientific theories — factors not recognized by realists who are anxious to demarcate their position from that of the anti-realists, and who are possessed of a disproportionate concern for the aspect of objectivity of truth, and its epistemic value in the description of nature.

The social beliefs, cognitive preferences, and modes of thought cannot, however, be in the nature of *determining* factors in the continuing unfolding of scientific theories describing and decoding the mechanisms inherent in deeper and deeper layers of reality. At the end of the day, the theories have to stand the test of reality and hence, whatever individual and social factors get involved in the formation and acceptance of hypotheses and theories, these are at best in the nature of *conditioning* factors.

Philip Kitcher's take on scientific progress is somewhat in the same spirit as Kuhn's:

“... I offer an analogy to show how my picture of science as providing objective knowledge does not entail that there is some unique, context-independent goal toward which inquiry aims. That analogy will also suggest a quite different way of thinking about the goals of the sciences and about scientific progress.” ([72]), Introduction).

The analogy Kitcher speaks of is that of map-making that illustrates the type of scientific realism he wants to espouse — the one he refers to as ‘modest realism’, where scientific progress is not driven teleologically to some distant but fixed epistemic goal, but is geared to meeting and solving problems of a contingent nature. It is this process that is socially conditioned. In particular, Kitcher relates the issue of scientific progress with that of social values.

The ideas about the scientific process raised in the present essay are in line with those developed by Kuhn and Kitcher.

Evolutionary psychology and cultural inheritance

All our cognitive quests are rooted in psychological processes in individuals and in social groups and communities. What is more, all acts of cognition are continuations of a prolonged *evolutionary process* — a process that is now recognized as forming an integral part of biological evolution — one that can be referred to as the process of *cognitive evolution* [51]. The idea of cognitive evolution may be broadened to include two aspects — a long term aspect within the process of biological evolution and a short term aspect of adaptation in the course of developmental history of individuals and societies, where the latter makes use of and realizes the capacities inherited in the evolutionary process. Evolutionary psychology talks of *dual inheritance* ([130], chapter 14) — modes of thought and cognitive traits acquired by genetic inheritance, and those acquired by *cultural inheritance*. Ideas, concepts, and modes of thought propagate culturally by processes that have a resemblance to the process of evolution by genetic means, though one need not read too much in this resemblance. Cultural inheritance has features of its own, and is seen to be expressed in various different cognitive traits in and across cultures distributed over the globe [92]. Inferences, hypotheses, and theories developed

by individuals and by scientific communities within different cultures are stamped by deep-rooted cognitive traits differing across cultures while, *at the same time*, possessing epistemic value. This is somewhat like the fact that competent chefs from different culinary cultures all cook excellent and delicious food, but all cook differently, and the food cooked by them have different *kinds* of taste. Or like the fact that great musicians from different cultural backgrounds all produce great and stirring music, but music of different *kinds*.

Evolutionary and social roots of human cognition, as also of cognitive values, have been insightfully discussed in [120], [121].

Basic ideas in evolutionary psychology can be found in [26], [17], [113], in addition to references mentioned earlier.

Cultural inheritance entails variations of modes of thought, beliefs, and value systems across cultures as also within cultures — since there exist *sub*-cultures within cultures (of which the family is an instance, being the bearer of a *micro*-culture) — all of which result in a conditioning of the interpretation of reality that, principally, is the business of science. All these taken together contrive to set the general course of scientific exploration that the individuals within a culture or a subculture undertake, the way of looking at problems and anomalies that come up, the relative importance of intuitive and deliberative cognitive resources brought to bear in the process of solving the problems, and a host of similar other features of the inferential and scientific process. It is, at times, naively asked as to how, for instance, the statement of Newton's laws might depend on cultural factors, since these leave too little room for cultural specificities to operate. The response to this cannot, of course, consist of posing an alternative statement of Newton's laws that a different cultural milieu could possibly produce, since cultural resources exert only broad structural features of scientific theory making. Copious indications of that are to be found in comparative studies of the sciences developing within continental, Islamic, Chinese, Indian, and similar other cultural formations during the later middle ages when the modes of practicing science were not subjugated to

one single mode of approach. Each of these cultures produced its own scientific framework committed to understanding some aspect or other of the mechanisms inherent in nature, and each having its own area of epistemic validity.

History proceeded along a course where most of the varied frameworks were all sought to be subsumed under a single dominating cognitive mode of doing science. It is no use wailing over how things could possibly have developed along an alternative course. But it is definitely a meaningful proposition to recognize the cross-cultural influences in science, to recognize that science is an interpretation of reality that is conditioned by social-cultural modes of thought, and to *make the best use* of this realization in the future development of the sciences.

The idea of differing cognitive traits across cultures — and *within* cultures too — has distinct dangers of misuse lurking underneath. What is a *difference* in the overall type or complexion, somewhat like a difference in flavor, and is often not amenable to explicit description, is in danger of being depicted as a matter of *superiority and inferiority* among culturally inherited cognitive traits. Such interpretations in terms of superiority and inferiority, and concomitant power relations of domination and subjugation invite, as a *reaction*, the opposite tendency of ignoring the cultural differences altogether, this notwithstanding the fact that cognitive differences have indeed been made use of in history in developing and extending power relations. This is where a deep and broad understanding of the notion of epistemic and non-epistemic aspects of the scientific process meets with troubled and turbulent waters.

The rationality of science: values in troubled times

Indeed, the scientific process as a whole is riddled with deeply conflicting aspects that seem as turbulent as the erupting conflicts in all major spheres of human interaction in today's world. This is because of the fact that the course of the scientific process is often, and indeed generally, dominated by *power interests*. The aim of scientific inquiry in major areas of investigation, considered from the instrumental point of view, is not neutral from the perspective of power relations. As for the epistemic aim, it is often of a

derivative nature, though a constantly present one, as we have seen above.

Even when all the non-epistemic aspects setting the course of the scientific process are ignored (this is a big counterfactual assumption indeed) the question of *rationality* of science is far from being a settled one. This is an area where, to start with, it appears that philosophical considerations of the twentieth century vintage (the first three-fourths of the century, that is,) can have a major role to play. But even such a limited discourse on the rationality of science does not inspire confidence, since the issue of rationality of science as a whole (assuming that the notion of science ‘as a whole’ is at all meaningful) turns out to be somewhat like that of the rationality of an individual, where the general feeling is that the question of rationality itself is not a well posed one. Indeed, in order to pronounce upon the rationality of the scientific process as a whole, one needs to know what the distinctive *method* of science is, and then to evaluate whether that method is conducive to achieving the epistemic aim (recall that we have agreed to ignore, like a true philosopher, all the non-epistemic aims) of science. And, the prospect is by no means encouraging on either of the two counts. Whether science has any distinctive method of its own, setting it apart from other spheres of human endeavor (the *demarcation* question in the philosophy of science; see, for background, [54]) is a question that does not appear to have a clear answer. That the concept of such a distinctive ‘method of science’ is not a sound one, has famously been expressed by Paul Feyerabend in his aphorism, ‘Anything goes’ ([38], p 19).

This widely quoted pithy statement of Feyerabends is, at times, interpreted out of context:

“It [the book ‘Against Method’] is not a systematic treatise; it is a letter to a friend and addresses his idiosyncrasies. For example, Imre Lakatos was a rationalist, hence rationalism plays a large role in the book. He also admired Popper and therefore Popper occurs much more frequently than his ‘objective importance’ would warrant. Imre Lakatos, somewhat jokingly, called me an anarchist and I had no objection to putting on the anarchist’s mask. Finally, Imre Lakatos loved to embarrass serious opponents with jokes and irony and so I, too, occasionally wrote in a rather ironical

vein. An example is the end of Chapter 1: ‘anything goes’ is not a ‘principle’ I hold — I do not think that ‘principles’ can be used and fruitfully discussed outside the concrete research situation they are supposed to affect — but the terrified exclamation of a rationalist [i.e., Lakatos] who takes a closer look at history. Reading the many thorough, serious, long-winded and thoroughly misguided criticisms I received after publication of the first English edition I often recalled my exchanges with Imre; how we would both have laughed had we been able to read these effusions together” ([38], preface).

The commonly accepted view of the scientific method is that science progresses in cycles of hypothesis (resulting in a theory), deduction, testing against evidence, and confirmation. Among these, the ‘theory of confirmation’ (either based on a principle of inductive generalization or of a more general probabilistic variety) has been seen to be not a solidly founded one and was replaced with a ‘theory of refutation’ by Karl Popper, which too did not prove to be beyond criticism. As for the process of hypothesis formation and the subsequent process of testing the deductive consequences against evidence, with these two phases of the scientific process forming, respectively, the context of discovery and the context of justification, the former was seldom addressed in the philosophy of science before the advent of the cognitive-naturalist era beginning, roughly, from the nineteen seventies, and the latter process (that of justification against evidence) was shown to be *theory-laden* (see, for instance, [12]), thus making the scheme of *theory* \rightarrow *deductive consequences* \rightarrow *justification against evidence* devoid of much of the rigor one would like it to have. In the resulting confusion, the very concept of a formal demarcation between science and non-science is all but abandoned though, paradoxically, there appears to be unanimity on the judgment about *particular* systems and practices, by adherents to diverse views in the philosophy of science ([54]; i.e., in other words, there have to be *some criteria* underlying our judgments in this respect that cannot, perhaps, be stated in formal terms).

At the same time, the notion of the epistemic aim of science, i.e., of scientific truths of nature, has also turned out to be a turbid one since, as we have seen, the concept of

'truth' is under a great state of stress between its non-epistemic origin and its epistemic responsibilities.

All in all, then, it seems that the issue of rationality of science is, again, a confused one.

But, at the same time, a criticism is often leveled against the cognitive-naturalist approach that the naturalist point of view itself is no less confusing since it is concerned with only a *descriptive* account of nature and natural mechanisms, including one of cognitive processes, and glosses over every concern with what *need be* or *ought to be*, thereby distancing itself with efforts at improvement of human cognitive limitations and, more broadly, of the human condition in general.

What is the position of naturalism on the question of *improvement*? This is an issue of stupendous magnitude, especially in today's world, where emotions, including destructive ones associated with greed, lust, and craving for power, have been let loose, and questions of enormous complexity are now flung out to riot-ridden streets, to be settled at dagger-point. Gone, perhaps, are the days of quiet and nuanced philosophical discourse, since momentous issues hang in the balance, to be decided under pressure of raw emotions at polling booths. Since this is what *is*, why torment ones mind with what *ought to be*? The cognitive-naturalist point of view appears to concern itself with describing what *is*, while remaining unconcerned with questions relating to improvement of what is, towards what ought to be.

It is indeed difficult to deny such accusations and adopt a stance of nonchalance, and it needs a serious engagement with real issues troubling the minds of men, where these issues are, in the main, not those of the mechanisms of cognition, but problems of searing intensity relating to *values* — to morality and ethics that have found no resolution in class-rooms, laboratories and philosophy conferences.

And the question of values ultimately hangs heavy on science itself. Science — and the fetishism of it — has been the cause of immense ravages, and that notwithstanding its supposed quest for truth, its concern with the secrets of nature. The real test of the

cognitive-naturalist approach will lie in finding an answer to how values and goals are generated in the world of individuals, communities, and cultures, and what role they are destined to play in the turbulent mental life that people live. This, of course, is vitally linked with how we perceive the world, and how we perceive *ourselves*.

The greatest philosophical traditions in history have time and again returned to this question: *how we perceive ourselves*. And the science of cognition itself is under no mean burden of responsibility in harnessing its efforts to this eternal quest of mankind — to that of understanding its own predicament in this world. This, of course, will entail a pressing engagement with greatly confusing questions of the self, of the soul, of volition and free will, of conscious and unconscious psychological processes in men, and of the deep existential turmoil inherent in the human condition of today's world.

Inductive inference appears to be just another piece of esoteric play of ideas, indifferent to the pressing concerns of our times and to the common sense perception of men. But inductive inference is precisely the mode of thought that defines the common sense perception: it is the typically exploratory mode of cognition that subsumes logical and philosophical thought but does not let itself be reduced to the latter. It is inductive inference through which we perceive the world and, in the process, perceive ourselves. But, whether an understanding of inductive inference with its associated unconscious cognitive processes will bring us any nearer to solving our own problems is anybody's guess, now that our world is decomposing and getting torn apart, not the least by the use of fruits of science itself.

Chapter 10

Summary and concluding words

Faltering ahead with a torch in hand

In this book we have taken a close look at two aspects of human inferential activity (see below), of which the making of scientific inferences constitutes a special instance. The two aspects are apparently contrary to each other, which makes the inferential process, or, to be more precise, our *understanding* of that process highly non-trivial and problematic.

It is to be clearly appreciated at the outset that inference making is the activity of the cognitive system aimed at making sense of the world and achieving certain ends, where the aspects of *achieving* and *making sense* have an intricate relation to each other. In these acts of making sense and achieving of ends, the cognitive system engages in an act of *interpreting* the world around it. Now, in order to understand how exactly the cognitive system interprets the world, we engage ourselves in understanding and interpreting the *activities of the cognitive system itself*. Thus, there are two levels of interpreting and making sense: one is the cognitive system making sense of the world at large, and the other is the system engaging in the reflective act of understanding its own inferential processes, especially those relating to the sciences. Of the two, the latter is actually a special instance of the former since the cognitive system is itself a part of the world at large, though that truism does not help us much except in realizing that both of these

are *open-ended* processes.

The open-endedness will get explained as I now recall the two apparently contrary aspects of the human inferential activity (including the one of making scientific inferences), which has been the recurrent theme in this book: *inferences are rooted in the cognitive system of the individual and of groups of individuals (such as the scientific communities in the case of scientific inferences) while, at the same time, these are aimed at grasping relevant aspects of a mind-independent reality.*

The inferential process of the individual is fundamentally of an inductive nature, where logically driven deductive processes find their place only *within* the ongoing inductive flow. The inductive process is generally in the nature of a deeply personal one, where beliefs and emotions play an essential role though, paradoxically, the end product of the inferential process has to be *effective* with reference to the world out there, a world that does not know of the personal beliefs and emotions. For instance, the scientific hypotheses and theories have to correctly correspond to features and mechanisms inherent in nature or, in other words, have to be bearers of *truth*.

The *personal* aspect of knowledge about the world was highlighted by Michael Polanyi who, moreover, stressed upon the *tacit* nature of the greater part of that knowledge and of the process of gaining that knowledge. Later generations of cognitive scientists worked upon the idea of tacit cognition, thereby bringing out a number of basic features of the cognitive activity of the human mind, including the *irrationality* of the cognitive process.

The 'irrationality' becomes apparent when individuals are given psychological 'tasks' of various types by way of requiring them to address little problems whose solutions are known to the test-givers (the cognitive psychologists) in terms of sets of *rules* supposed to be *normative* ones. It then becomes necessary to recognize the distinction between the normative and the *descriptive*, the latter being the way the cognitive and inferential processes *actually* proceed within the human mind. Now, this is a tricky question that requires a deep look at how the cognitive process operates tacitly, i.e., at an *unconscious*

level.

The unconscious mind is capable of greatly complex cognitive activities, most of which were previously assumed to be exclusively dependent on focused awareness and intention. While some revealing indications of the range of activities of the cognitive unconscious are available by means of psychological studies, there remains a huge unexplored area, which calls for speculations and interpretations so that the few pieces of solid information available can be interpreted and woven into a coherent framework. Evidently, there may be alternative frameworks of interpretation and speculation, some of which will prove to be inadequate in the light of subsequent findings, while some others will gain in strength, and this process of speculating, pruning, and gaining in strength will continue and, additionally, there will be occasional *major* transformations in our conception of how the cognitive process, including the processes of inference, works.

This, incidentally, is *also* the way that science approaches nature and natural phenomena. Mechanisms operating within natural phenomena are revealed to science in successive stages, in each of which nature appears in a new perspective as greater and greater depths are probed by means of improved instruments, backed up with improved conceptual frameworks. In other words, at every stage of interpretation of reality, there exist un-examined substrata and unacknowledged natural boundaries, due to which the fit between nature and the theories of science shows up small anomalies and faults symptomatic of an innate vulnerability of these theories. The latter get replaced by broader and deeper theories at a later stage which, however, do not bear a clear and simple relation of reducibility to the earlier generation of theories. The transition from one stage to the next is often a complex process where hypotheses, essentially of the nature of inspired guesses, are proposed, and highly speculative ideas are subjected to criticisms, counter-criticisms, and meticulous appraisal, examining their consistency with evidence and with vast bodies of concepts of proven worth. In the process, some of the speculations and hypotheses get abandoned while some others are selected for further scrutiny and examination.

In the special case of building of theories aimed at describing our cognitive and inferen-

tial processes, the vulnerability of speculations, hypotheses, and theories is especially pronounced because of the fact that most of these processes are in the nature of hidden ones, taking place within the cognitive unconscious. However, one seems to be on fairly solid ground in saying that all the inferential processes are contextually determined, where there is an *external* context as also an *internal* one to be reckoned with, both of which are of a deeply complex nature. The external context, described by the stimuli and inputs from the external world that sets an inferential process in motion is complex in virtue of the fact that out of the infinite number of environmental inputs that the cognitive mind can possibly pick up, only a subset (which may be *a very large one*) is actually selected depending on their relevance and salience. However, the selection is made complex by the fact that, first, the criteria of relevance and salience depends on the *current state of mind* (including the current set of goals, purposes and values) of the individual concerned and, secondly, many of the inputs received from the environment, and subsequently used for inferential purposes, are of a *subliminal* nature. Most of the inputs are in the form of cues picked up tacitly and subsequently transformed into unconsciously formed heuristics.

The internal context of an inferential process is of a similarly complex nature, if not more, since it is made up of heuristics and beliefs, mostly of an unconscious nature, *additionally* involving emotions and affects. In other words, there exist vast repertoires of *hidden* components in both the external and internal contexts relevant to the inferential process.

An inference involves a processing of information of a tangled and complex nature where a vast hierarchy of *rules* are made use of, again, mostly at an unconscious level. The rules are of various categories, some of which are independent of beliefs and modes of thought of individuals and groups of individuals while some others are *not so*. The latter include rules of a person-specific and those of an inter-subjective nature, mostly answering to the description of *heuristics*. The latter are half-baked rules of thumb and hunches, many of those of a transient and fluid nature, liable to be discarded unless proven to be of some worth. On the other hand, there is a vast web of beliefs, many of which are resistant to revision and are of a relatively remote relevance with reference

to the inferential act in question. What is of special significance is that a number of heuristics and beliefs act as resources that propel the cognitive mind across logical gaps, where the latter are gaps that cannot be bridged by means of objectively defined rules. In this, the cognitive mind is greatly aided by *emotions* and affects that facilitate the making of decisions and acquire relevance in the bridging of the logical gaps. It is this process of leaping across logical gaps that is specific to *inductive* inference, which makes it fallible and, at the same time, *personal* in nature while being uniquely *effective* too.

Of special relevance in the sciences is the process of abduction, i.e., the one of making of *hypotheses* that subsequently germinate into scientific theories. A novel hypothesis that leads to the transformation of an entire conceptual space is a mysterious process indeed, the true nature of which can only be speculated upon. It is likely to involve a greatly enhanced exploration of the conceptual space where local instabilities in the sequential progression of information processing, caused by the amplifying action of emotions, lead to the development of parallel branchings in the exploration process. In this, efficient organizing principles like the detection of analogies play a crucial role whereby relatively remote ideas get correlated and eventually coalesce together to lead to a conceptual transformation engendering remarkable possibilities.

The scientific theories that result from these processes of an abstruse nature have a dual significance: these sprout from grounds abounding in beliefs and cultural resources of individuals and groups of individuals, and are built upon prior structures of *existing* concepts and theories and, *at the same time*, these are aimed at revealing the inner mechanisms of nature. As a result, the theories are bearers of *truths* about nature that have strange and conflicting aspects in them. On the one hand, they are truths largely independent of opinions and points of view of individuals and groups of individuals (though this is conditional upon a more or less prolonged process of exchanges and communications of ideas) and, in this sense, are *objective* in nature while, on the other, they are aimed at a mind-independent reality. The mechanisms inherent in the latter are explored in successive stages of theory that are of an incommensurable nature, through conceptual transformations akin to a change in perspective. In this, the truths

are in the nature of *socially conditioned interpretations*.

But the fact that the inferential acts in science, including the ones arrived at by the process of abduction, are fundamentally in the nature of interpretations, with logical gaps remaining within, need not imply a weakness in these. On the contrary, it constitutes a *strength* of great relevance in our inquiry into nature and in our incessant engagements with a largely unknown, uncertain, and complex world. There does not exist any sure-shot way of coming to grips with a vast and complex reality with our limited and meager cognitive resources other than the one of guessing and groping for our way ahead, sticking our neck out, making use of what has proved to be of some worth and discarding what is found to be ineffective while, at the same time, retaining the lessons of the failures.

The cognitive abilities are, to a quite considerable extent, results of a protracted evolutionary process that is essentially of a similar nature as the one outlined above — building upon the past in a piecemeal way, in response to *contingent* necessities. Cognition has no ultimate goal precisely because it faces nature in an infinite variety of contexts. Inferences, abstractions, and theories are not aimed at producing a facsimile description of how nature exists and behaves as a whole, but at providing us with a *summary understanding* of parts and aspects of nature as we face these in specific contexts so that we can make *effective sense* of these. Our theories are like *maps* drawn from a *limited* exploration into nature, based on which we make hypotheses regarding the way ahead. We then commit ourselves to further explorations in keeping with our hypotheses and, when rewarded with success, remake the maps, where the new set of maps differ from the earlier ones in that new *aspects* are incorporated into these, requiring novel ways of reading the maps. Theories, in other words, make possible new *encounters* with nature. In this, science is continuous with and constitutes a heightened form of our mundane, day-to-day engagement with reality where we make great use of our ability to guess, and guess correctly, albeit with equally great support from our judgment based on sound logic.

The naturalist point of view looks at the actual process by which scientists go about

their business of interpreting nature, without burdening itself with *logical* and abstract considerations of what the aim of science is, what the scientific method is, and what constitutes scientific progress. In this, the naturalist point of view focuses keenly on the *cognitive roots* of how science inquires into nature. It focuses keenly on how inferential processes actually proceed in the minds of individuals, how a great variety of cultural resources affect our cognitive endeavor, and how hypotheses are actually formed, giving rise to theories about the world around us, without burdening itself too much with questions of *norms* of rationality; more precisely, it tells us that questions relating to norms are, in a sense, misplaced ones.

Still, questions relating to norm are *not irrelevant*. And, abstract and logical considerations are not irrelevant either. In a manner of speaking, considerations *in the abstract* are as relevant as those *in the concrete*. In the context of human cognition and of the scientific process, the naturalist point of view entails the considerations of the latter variety, while the former are the ones that were the mainstay in the philosophy of science up to the sixties of the last century. The logical approach in the understanding of human cognition, principally geared to realizations of cognitive mechanisms in artificial intelligence, captures quite an impressive number of aspects of the cognitive process, and a logical analysis of the scientific process brings in sharp focus quite an impressive number of issues relating to the cognitive roots of scientific exploration and the way the latter relates to a mind-independent reality.

Questions of norm in human cognition cannot be shaken off by the simple assertion that these do not relate to how cognition *actually* works. Human cognitive and inferential processes are not limited within the narrow horizon of effectiveness, because effectiveness is meaningful only with reference to *goals* and *values*. And, questions relating to goals and values are not confined to the field of cognition alone, because these are deeply existential ones. Ultimately, these relate to our desires and drives, our cravings for power, our yearnings for fulfilment, our deep-rooted instincts for sharing and understanding, our need to *improve* upon what we have become, and our endless quest for *making sense* of our own existence in this world of ours a world that is *within* us as it is *around* us.

Both the two worlds are fathomless, revealed to us contextually and discontinuously, in bits and pieces. Our perception of either of the two builds up, shimmers before our eyes, and then dissolves into a new picture, revealing novel aspects in a new context. This makes for a quest that remains open-ended even as it constitutes an intoxicating and dizzying journey. We do hold the powerful torch of *logic* in our hands, but the rays emanating from it are too straight to obviate the necessity of guessing, groping, interpreting, and faltering ahead along twisting paths in a world that is at once labyrinthine and layered.

Bibliography

- [1] Frederic Perez-Alvarez, and Carme Timoneda-Gallart, *A Better Look at Intelligent Behavior: Cognition and Emotion*, Nova Science Publishers, Inc., New York, (2007).
- [2] David M. Armstrong, *Belief, Truth and Knowledge*, Cambridge University Press, Cambridge (1973).
- [3] Elaine Perry, Daniel Collerton, Fiona LeBeau, and Heather Ashton, (ed.) *New Horizons in the Neuroscience of Consciousness*, John Benjamins Publishing Company, Amsterdam (2010).
- [4] Ellaine Perry, Heather Ashton, and Allan Young, (ed.) *Neurochemistry of Consciousness*, John Benjamins Publishing Company, Amsterdam (2002).
- [5] Lynne R. Baker, *Saving Belief: A Critique of Physicalism*, Princeton University Press, Princeton (1987).
- [6] Robert W. Batterman, *The Devil in the Details: Asymptotic Reasoning in Explanation, Reduction, and Emergence*, Oxford University Press, Oxford (2002).
- [7] Johan van Benthem, *Modal Logic for Open Minds*, Center for the Study of Language and Information, Stanford University, Stanford (2010).
- [8] Morten L. Kringelbach and Kent C. Berridge, (ed.) *Pleasures of the Brain*, Oxford University Press, Oxford (2010).
- [9] Michael V. Berry, 'Asymptotics, Singularities, and the Reduction of Theories', in: *Logic, Methodology and Philosophy of Science IX*, Dag Prawitz, Brian Skyrms, and Dag Westerstahl (eds.), Elsevier, Amsterdam, section 11, pp 597-608 (1994).

BIBLIOGRAPHY

- [10] Benjamin M. Bly, and David E. Rumelhart (eds.), *Cognitive Science*, Academic Press, San Diego (1999).
- [11] Margaret A. Boden, *The Creative Mind: Myths and Mechanisms*, Routledge, London (2004).
- [12] James Bogen, "Theory and Observation in Science", The Stanford Encyclopedia of Philosophy (Summer 2017 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/sum2017/entries/science-theory-observation/> (2017).
- [13] Alisa Bokulich, *Reexamining the Quantum-Classical Relation: Beyond Reductionism and Pluralism*, Cambridge University Press, Cambridge (2008).
- [14] William F. Brewer and Bruce L. Lambert, 'The theory-ladenness of observations and the theory-ladenness of the rest of the scientific process', *Philosophy of Science*, 68 (S3): S176-S186 (2001).
- [15] Charles Dunbar Broad, *Ethics and the History of Philosophy*, Humanities Press, New York (1952).
- [16] Wilma Bucci, 'The Need for a Psychoanalytic Psychology in the Cognitive Science Field', *Psychoanalytic Psychology*, 17(2):203-224 (2000).
- [17] David M. Buss (ed.), *The Handbook of Evolutionary Psychology, vol.1, Foundations*, John Wiley & Sons, Inc., Hoboken (2016).
- [18] Nancy Cartwright, *How the Laws of Physics Lie*, Oxford University Press, Oxford (1983).
- [19] Sergio Chibbaro, Lamberto Rondoni, and Angelo Vulpiani, *Reductionism, Emergence, and Levels of Reality: The Importance of Being Borderline*, Springer, Switzerland (2014).
- [20] Paul M. Churchland, *Scientific Realism and the Plasticity of Mind*, Cambridge University Press, Cambridge (1979).

BIBLIOGRAPHY

- [21] Axel Cleeremans, *Mechanisms of Implicit Learning: Connectionist Models of Sequence processing*, a Bradford Book, MIT Press, Massachusetts (1993).
- [22] Axel Cleeremans, 'Connecting Conscious and Unconscious Processing', *Cognitive Science*, 38:. 1286-1315 (2014).
- [23] Harry Collins, *Tacit and Explicit Knowledge*, The University of Chicago Press, Chicago (2010).
- [24] Michael H. Connors and Peter W. Halligan, 'A cognitive account of belief: a tentative road map', *Front. Psychol.*, 5:1588-1607 (2014).
- [25] Roberta Corvi, *An Introduction to the Thought of Karl Popper (Translated by Patrick Camiller)*, Routledge, London and New York (1993).
- [26] Leda Cosmides and John Tooby, 'Better than Rational: Evolutionary Psychology and the Invisible hand', *American Economic Review*, 84(2):327-332(1994).
- [27] Nicole Van Hoeck, Patrick D. Watson, and Aron K. Barbey, 'Cognitive neuroscience of human counterfactual reasoning', *Front. Hum. Neurosci.*, 9:420 (2015).
- [28] Vincenzo Crupi, 'Confirmation', *The Stanford Encyclopedia of Philosophy* (Winter 2016 Edition), Edward N. Zalta(ed.), URL = <https://plato.stanford.edu/archives/win2016/entries/confirmation/> (2016).
- [29] Denise D. Cummins, 'The evolution of Reasoning', in: *The Nature of Reasoning*, Jacqueline P. Leighton and Robert J. Sternberg (eds.), Cambridge University Press, Cambridge, chapter 13 (2004).
- [30] Antonio R. Damasio, *Descartes' Error: Emotion, Reason, and the Human Brain*, Avon Books, New York (1994).
- [31] Antonio R. Damasio, *Self Comes to Mind: Constructing the Conscious Brain*, Pantheon Books, New York (2010).
- [32] William J. Devlin and Alisa Bokulich (eds.), *Kuhns Structure of Scientific Revolutions — 50 Years On*, Springer, Switzerland (2015).

BIBLIOGRAPHY

- [33] Antonio Diguez, 'Kitchers Modest Realism: The Reconceptualization of Scientific Objectivity', in: *Scientific Realism and Democratic Society: The Philosophy of Philip Kitcher*, Wenceslao J. Gonzalez (ed.), Brill, Leiden (2011).
- [34] Ronald de Sousa, 'Epistemic Feelings', in: *Epistemology and Emotions*, Georg Brun, Ulvi Doguoglu, and Dominique Kuenzle, (eds.), Ashgate Publishing Limited, Hampshire, chapter 9 (2008).
- [35] Eric Eich, John F. Kihlstrom, Gordon H. Bower, Joseph P. Forgas, and Paula M. Niedenthal, (eds.), *Cognition and Emotion*. Oxford University Press, Oxford (2000).
- [36] Rene Elio (ed.), *Common Sense, Reasoning, and Rationality*, Oxford University press, Oxford (2002).
- [37] Jonathan St. B. T. Evans and David E. Over, *Rationality in Reasoning*, Psychology Press, Sussex (1996).
- [38] Paul Feyerabend, *Against method*, Verso, London (1993).
- [39] Klaus Fiedler and Herbert Bless, 'The formation of beliefs at the interface of affective and cognitive processes', in: *Emotions and Beliefs: How feelings Influence Thoughts*, Nico H. Frijda, Antony S. R. Manstead, and Sacha Bem (eds.), Cambridge University Press, Cambridge, chapter 6 (2000).
- [40] Peter A. Flach and Antonis C. Kakas, *Abduction and Induction*, Springer Science+Business Media B.V, Dordrecht (2000).
- [41] Nico H. Frijda and Batja Mesquita, 'Belief through emotions', in: *Emotions and Beliefs: How feelings Influence Thoughts*, Nico H. Frijda, Antony S. R. Manstead, and Sacha Bem (eds.), Cambridge University Press, Cambridge. chapter 3 (2000).
- [42] Peter Grdenfors, *Conceptual spaces: the geometry of thought*, MIT press, Cambridge (1991).
- [43] Ronald N. Giere, *Scientific Perspectivism*, The University of Chicago Press, Chicago (2006).

BIBLIOGRAPHY

- [44] Gerd Gigerenzer, *Gut Feelings: The Intelligence of the Unconscious*, Viking, Penguin Group (USA) Inc. (2007).
- [45] Gerd Gigerenzer and Reinhard Selten (eds.), *Bounded rationality: The Adaptive Toolbox*, MIT Press, Massachusetts (2002).
- [46] Peter Godfrey-Smith, *Theory and Reality*, The University of Chicago Press, Chicago (2003).
- [47] Vinod Goel and Oshin Vartanian, 'Negative emotions can attenuate the influence of beliefs on logical reasoning', *Cognition and Emotion*, 25(1):121-131 (2011).
- [48] Wenceslao. J. Gonzalez, 'From Mathematics to Social Concern about Science: Kitcher's Philosophical Approach', in: *Scientific Realism and Democratic Society: The Philosophy of Philip Kitcher*, Wenceslao. J. Gonzalez (ed.), Brill, Leiden (2011).
- [49] Wenceslao J. Gonzalez (ed.), *Scientific Realism and Democratic Society: The Philosophy of Philip Kitcher*, Brill, Leiden (2011).
- [50] Ian Hacking, *Representing and Intervening*, Cambridge University Press, Cambridge (1983).
- [51] Edward H. Hagen, 'Evolutionary Psychology and its Critics', in: *The Handbook of Evolutionary Psychology*, vol. 1, Foundations, David M. Buss (ed.), John Wiley and Sons, Inc., Hoboken, chapter 4 (2016).
- [52] Norwood R. Hanson, *Patterns of Discovery*, University of Cambridge Press, Cambridge (1958).
- [53] Norwood R. Hanson, 'Is There a logic of Scientific Discovery?', *The Australian Journal of Philosophy*, 38(2):91-106 (1960).
- [54] Sven Ove Hansson, 'Science and Pseudo-Science', *The Stanford Encyclopedia of Philosophy* (Summer 2017 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/sum2017/entries/pseudo-science/> (2017).
- [55] Ran R. Hassin, 'Nonconscious Control and Implicit Working Memory', in: *The New Unconscious*, Ran R. Hassin, James S. Uleman, John A. Bargh, (eds.), Oxford University Press, Oxford, chapter 8 (2005).

BIBLIOGRAPHY

- [56] Ran R. Hassin, James S. Uleman, John A. Bargh, (eds.), *The New Unconscious*, Oxford University Press, Oxford (2005).
- [57] Evan Heit, 'What Is Induction and Why Study It?', in: *Inductive Reasoning: Experimental, Developmental, and Computational Approaches*, Aidan Feeney and Evan Heit (eds.), Cambridge University Press, Cambridge, chapter 1 (2007).
- [58] Carl G. Hempel and Paul Oppenheim, 'Studies in the Logic of Explanation', *Philosophy of Science*, 15(2):135-175 (1948).
- [59] Claire Hewson, 'Empirical Evidence Regarding the Folk Psychological Concept of Belief', URL=www.academia.edu/4484057/Empirical_evidence_regarding_the_folk_psychological. [Last visited: 30.11.2017].
- [60] Douglas Hofstadter and Emmanuel Sander, *Surfaces and Essences: Analogy as the Fuel and Fire of Thinking*, Basic Books, New York (2013).
- [61] John H. Holland, Keith J. Holyoak, Richard E. Nisbett, and Paul R. Thagard, *Induction: Processes of Inference, Learning, and Discovery*, Massachusetts Institute of Technology, Massachusetts (1993).
- [62] Keith J. Holyoak and Paul Thagard, *Mental Leaps: Analogy in Creative Thought*, A Bradford Book, The MIT Press, Cambridge (1995).
- [63] Jan De Houwer and Dirk Hermans, *Cognition & Emotion: Review of Current Research and Theories*, Psychology Press, New York (2010).
- [64] G. Randolph Mayes, 'Theories of Explanation', Internet Encyclopedia of Philosophy, URL=<http://www.iep.utm.edu/explanat/#SH4e> [Last visited: 04.12.2017]
- [65] Hilde M. Huizenga, Ruud Wetzels, Don van Ravenzwaaij, Eric-Jan Wagenmakers, 'Four empirical tests of Unconscious Thought Theory', *Organizational Behavior and Human Decision Processes*, 117: 332-340 (2012).
- [66] Struan Jacobs, 'Polanyis presagement of the incommensurability concept', *Studies in History and Philosophy of Science*, 33:105-120 (2002).

BIBLIOGRAPHY

- [67] Stefania R. Jha, 'Wigners "Polanyian" Epistemology and the Measurement Problem: The Wigner-Polanyi Dialog on Tacit Knowledge', *Phys. Perspect.*, 13:329-358 (2011).
- [68] Philip N. Johnson-Laird, *How We Reason*, Oxford University Press, Oxford (2008).
- [69] Philip N. Johnson-Laird and Eldar Shafir, 'The interaction between reasoning and decision making: an introduction', *Cognition*, 49:1-9 (1993).
- [70] John F. Kihlstrom, 'The Cognitive Unconscious', *Science*, 237:1445-1452 (1987).
- [71] Philip Kitcher, *The Advancement of Science: Science without Legend, Objectivity without Illusion*, Oxford University Press, Oxford (1993).
- [72] Philip Kitcher, *Science, Truth, and Democracy*, Oxford University Press, Oxford (2001).
- [73] Robert Klee, *Introduction to the Philosophy of Science: Cutting nature at Its Seems*, Oxford University Press, Nw York, Oxford (1997).
- [74] Christopher Koch, 'Consciousness Redux: Probing the Unconscious Mind', *Scientific American*, Nov.-Dec. (2011)
- [75] Saul A. Kripke, *Wittgenstein on Rules and Private Language: An Elementary Exposition*, Harvard University Press, Cambridge, Massachusetts (1982).
- [76] Thomas S. Kuhn, *The Structure of Scientific Revolutions* (Second Edition, Enlarged), University of Chicago Press, Chicago, (1970).
- [77] Thomas Kuhn, *The Copernican revolution*, Harvard University Press (reprint), Cambridge, MA(1985).
- [78] James Ladyman, *Understanding Philosophy of Science*, Routledge, London (2002).
- [79] Mary C. Lamia, 'Feeling is Believing: Your emotional past biases our present beliefs', URL=<https://www.psychologytoday.com/blog/intense-emotions-and-strong-feelings/201210/feeling-is-believing>; [Last visited: 01.12.2017].
- [80] Jonathan Lear, *Freud*, Routledge, Abingdon (2005).

BIBLIOGRAPHY

- [81] Jacqueline P. Leighton, 'The Assessment of Logical Reasoning', in: *The Nature of Reasoning*, Jacqueline P. Leighton and Robert J. Sternberg (eds.), Cambridge University Press, Cambridge, chapter 11 (2004).
- [82] Lorenzo Magnani, *Abduction, Reason, and Science: Process of Discovery and Explanation*, Kluwer Academic / Plenum Publishers (2000).
- [83] Lorenzo Magnani, *The Abductive Structure of Scientific Creativity: An essay on the Ecology of Cognition*, Springer International Publishing AG, Switzerland (2017).
- [84] Mandelbaum, Eric, "Associationist Theories of Thought", The Stanford Encyclopedia of Philosophy (Summer 2017 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/sum2017/entries/associationist-thought/>.
- [85] Ken Manktelow, *Reasoning and Thinking*, Psychology Press, Sussex (1999).
- [86] Barnaby Marsh, Peter M. Todd, Gerd Gigerenzer, 'Cognitive Heuristics, Reasoning the Fast and Frugal Way', in: *The Nature of Reasoning*, Jacqueline P. Leighton and Robert J. Sternberg (eds.), Cambridge University Press, Cambridge, chapter 10 (2004).
- [87] James L. McClelland, and David E Rumelhart, *Parallel Distributed Processing. Explorations in the Microstructure of Cognition. Volume 2: Psychological and Biological Models* MIT Press, Cambridge, MA (1986).
- [88] Peter B. Medawar, *Induction and Intuition in Scientific Thought*, Methuen & Co Ltd, London (1970).
- [89] Brian Kolbe and Ian Q. Whishaw, *Fundamentals of Human Neuropsychology*, seventh edition, Worth Publishers, New York (2015).
- [90] Ikka Niiniluoto, 'Scientific Progress', The Stanford Encyclopedia of Philosophy (Summer 2015 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/sum2015/entries/scientific-progress/> (2015).

BIBLIOGRAPHY

- [91] Richard E. Nisbett, *Mindware: Tools for Smart Thinking*, Farrar, Straus and Giroux, New York (2015).
- [92] Richard E. Nisbett, *The Geography of Thought: How Asians and Westerners Think Differently and Why*, The Free Press, New York (2003).
- [93] A. Poddiakov (2013). 'Catching a flying ball — is that really that easy? A contribution to the critique of G. Gigerenzers approach', in: A.G.Egorov, V.V.Selivanov (eds.), *Psychologia kognitivnykh processov: Materialy 4 vserossiiskoi konferentsii* [Psychology of cognitive processes: Proceedings of the 4th All-Russian conference] (pp. 83-86; translated from Russian). Smolensk: Smolensk State University; URL=https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2484001 2013[Last visited: 06.01.2018].
- [94] Michael Polanyi, *Personal Knowledge: Towards a Post-Critical Philosophy*, Routledge & Kegan Paul Ltd, London (1962).
- [95] Michael Polanyi, *The Tacit Dimension* (with a new foreward by Amartya Sen, 2009), The University of Chicago Press Ltd., London (1966).
- [96] Michael Polanyi, *Science, Faith and Society*, Oxford University Press, London (1946).
- [97] Michael Polanyi, 'Science and Reality', *British Journal for the Philosophy of Science*, 18(3):177-196 (1967).
- [98] George Polya, *Induction and Analogy in Mathematics*, Princeton University Press, Princeton (1954).
- [99] William Poundstone, *Labyrinths of Reason*, Penguin Books, London (1988).
- [100] Mick Power and Tim Dalgleish, *Cognition and Emotion: From order to Disorder*, Psychology Press, New York (2008).
- [101] Hilary Putnam, *Mathematics, Matter and Method*, Cambridge University Press, Cambridge, MA (1975).

BIBLIOGRAPHY

- [102] Willard V. O. Quine and J. S. Ullian, *The Web of Belief*, McGraw-Hill, Inc., New York (1978).
- [103] Arthur Reber, *Implicit Learning and Tacit Knowledge*, Oxford University Press, New York, Oxford (1993).
- [104] Maxwell J. Roberts, 'Heuristics and Reasoning I: Making Deduction Simple', in: *The Nature of Reasoning*, Jacqueline P. Leighton and Robert J. Sternberg (eds.), Cambridge University Press, Cambridge, chapter 9 (2004).
- [105] Alex Rosenberg, *Philosophy of Science: A Contemporary Introduction*, Routledge, New York (2005).
- [106] Albert Rothenberg, *Flight from Wonder: An Investigation of Scientific Creativity*, Oxford University Press, Oxford (2015).
- [107] Oliver Sacks, and Joy Hirsch, 'A Neurology of Belief', editorial in *Annals of Neurology*, 63(2):129-130 (2008).
- [108] Jutta Schickore and Friedrich Steinle, *Revisiting Discovery and Justification: Historical and philosophical perspectives on the context distinction*, Springer, Dordrecht (2006).
- [109] Gerhard Schurz, *Philosophy of Science: A Unified Approach*, Routledge, New York (2014).
- [110] Herbert A. Simon, 'The Scientist as Problem Solver', in: *Complex Information Processing: The Impact of Herbert A. Simon*, David Klahr and Kenneth Kotovsky (eds.), Lawrence Erlbaum Associates, Inc., Hillsdale, NJ, chapter 14 (1989).
- [111] Lina Soler, Howard Sankey, and Paul Hoyningen-Huene (eds.), *Rethinking Scientific Change and Theory Comparison: Stabilities, Ruptures, Incommensurabilities?*, Springer, Dordrecht (2008).
- [112] Keith E. Stanovich, *Decision Making and Rationality in the Modern World*, Oxford University press, Oxford (2003).

BIBLIOGRAPHY

- [113] Kim Sterenly, *The Evolution of Human Cognition*, Blackwell Publishing, Malden, MA (2003).
- [114] Stephen P. Stich, *From folk psychology to cognitive science: The case against belief*, The MIT Press, Cambridge, MA (1983).
- [115] David Soto, Teemu Mntil, and Juha Silvanto, 'Working memory without consciousness', *Current Biology*, 21(22): R912-R913 (2011).
- [116] William Talbott, 'Bayesian Epistemology', *The Stanford Encyclopedia of Philosophy* (Winter 2016 Edition), Edward N. Zalta(ed.), URL = <https://plato.stanford.edu/archives/win2016/entries/epistemology-bayesian/> (2016).
- [117] Joshua B. Tenenbaum, Charles Kemp, and Patrick Shafto, 'Theory-Based Bayesian Models of Inductive Reasoning', in: *Inductive Reasoning: Experimental, Developmental, and Computational Approaches*, Aidan Feeney and Evan Heit (eds.), Cambridge University Press, Cambridge, chapter 7 (2007).
- [118] Paul Thagard, *Hot Thought: Mechanisms and Applications of Emotional Cognition*, A Bradford Book, The MIT Press, Cambridge, MA (20006).
- [119] Adam Timmins, 'Why Was Kuhns 'Structure' More Successful than Polanyi's 'Personal Knowledge?' ', *HOPOS: The Journal of the International Society for the History of Philosophy of Science*, 3(2):306-317 (2013).
- [120] Michael Tomasello, *A Natural History of Human Thinking*, Harvard University Press, Cambridge, MA (2014).
- [121] Michael Tomasello, *A Natural History of Human Morality*, Harvard University Press, Cambridge, MA (2016).
- [122] James S. Uleman, 'Becoming Aware of the New Unconscious', in: *The New Unconscious*, Ran R. Hassin, James S. Uleman, John A. Bargh, (eds.), Oxford University Press, Oxford, Introductory chapter (2005).
- [123] Bas van Fraassen, *The Scientific Image*, Oxford University Press, Oxford (1980).

- [124] Daniel M. Wegner, 'Who is the Controller of Controlled processes?', in: *The New Unconscious*, Ran R. Hassin, James S. Uleman, John A. Bargh, (eds.), Oxford University Press, Oxford, chapter 1 (2005).
- [125] Robert W. Weisberg and Laretta M. Reeves, *Cognition: From Memory to Creativity*, John Wiley & Sons, Inc., Hoboken, New Jersey (2013).
- [126] Drew Westen, 'The Cognitive Self and the Psychoanalytic Self: Can We Put Our Selves Together?', *Psychological Inquiry*, 3(1):1-13 (1993)
- [127] Eugene P. Wigner, 'The Unreasonable Effectiveness of Mathematics in the Natural Sciences', Richard Courant Lecture in Mathematical Sciences, New York University, May 11, 1959; *Comm Pure Appl. Math.*, 13:1-14 (1960).
- [128] Eugene P. Wigner and R. A. Hodgkin, 'Michael Polanyi. 12 March 1891 — 22 February 1976', *Biographical Memoirs of Fellows of the Royal Society*, 23(Nov., 1977):413-448 (1977).
- [129] Woodward, James, "Scientific Explanation", *The Stanford Encyclopedia of Philosophy* (Fall 2017 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/fall2017/entries/scientific-explanation/> (2017).
- [130] Lance Workman and Will Reader, *Evolutionary Psychology: An Introduction*, Cambridge University Press, Cambridge (2014).